



Evolutionary Psychology and the Emotions

Leda Cosmides & John Tooby

To appear in:

Handbook of Emotions, 2nd Edition
M. Lewis & J. M. Haviland-Jones, Editors.
NY: Guilford.

Copyright 2000 Leda Cosmides and John Tooby

Caution: The print version may differ in minor respects from this draft.
Posted only for scholarly/educational use. Please contact the publisher directly for permission to reprint.

Evolutionary psychology is an approach to the psychological sciences in which principles and results drawn from evolutionary biology, cognitive science, anthropology, and neuroscience are integrated with the rest of psychology in order to map human nature. By human nature, evolutionary psychologists mean the evolved, reliably developing, species-typical computational and neural architecture of the human mind and brain. According to this view, the functional components that comprise this architecture were designed by natural selection to solve adaptive problems faced by our hunter-gatherer ancestors, and to regulate behavior so that these adaptive problems were successfully addressed (for discussion, see Cosmides & Tooby, 1987, Tooby & Cosmides, 1992). Evolutionary psychology is not a specific subfield of psychology, such as the study of vision, reasoning, or social behavior. It is a *way of thinking* about psychology that can be applied to any topic within it - including the emotions.

The analysis of adaptive problems that arose ancestrally has led evolutionary psychologists to apply the concepts and methods of the cognitive sciences to scores of topics that are relevant to the study of emotion, such as the cognitive processes that govern cooperation, sexual attraction, jealousy, aggression, parental love, friendship, romantic love, the aesthetics of landscape preferences, coalitional aggression, incest avoidance, disgust, predator avoidance, kinship, and family relations (for reviews, see Barkow, Cosmides, & Tooby, 1992; Crawford & Krebs, 1998; Daly & Wilson, 1988; Pinker, 1997).

Indeed, a rich theory of the emotions naturally emerges out of the core principles of evolutionary psychology (Tooby 1985; Tooby & Cosmides, 1990a; see also Nesse, 1991). In this chapter we will (1) briefly state what we think emotions are and what adaptive problem they were designed to solve; (2) explain the evolutionary and cognitive principles that led us to this view; and (3) using this background, explicate in a more detailed way the design of emotion programs and the states they create.

An evolutionary psychological theory of the emotions

An evolutionary perspective leads one to view the mind as a crowded zoo of evolved, domain-specific programs. Each is functionally specialized for solving a different adaptive problem that arose during hominid evolutionary history, such as face recognition, foraging, mate choice, heart rate regulation, sleep management, or predator vigilance, and each is activated by a different set of cues from the environment. But the existence of all these microprograms itself creates an adaptive problem: Programs that are individually designed to solve specific adaptive problems could, if simultaneously activated, deliver outputs that conflict with one another, interfering with or nullifying each other's functional products. For example, sleep and flight from a predator require mutually inconsistent actions, computations, and physiological states. It is difficult to sleep when your heart and mind are racing with fear, and this is no accident: disastrous consequences would ensue if proprioceptive cues were activating sleep programs at the same time that the sight of a stalking lion was activating ones designed for predator evasion. To avoid such

consequences, the mind must be equipped with superordinate programs that override some programs when others are activated (e.g., a program that deactivates sleep programs when predator evasion subroutines are activated). Furthermore, many adaptive problems are best solved by the simultaneous activation of many different *components* of the cognitive architecture, such that each component assumes one of several alternative states (e.g., predator avoidance may require simultaneous shifts in both heart rate and auditory acuity; see below). Again, a superordinate program is needed that coordinates these components, snapping each into the right configuration at the right time.

Emotions are such programs. To behave functionally according to evolutionary standards, the mind's many subprograms need to be orchestrated so that their joint product at any given time is functionally coordinated, rather than cacophonous and self-defeating. This coordination is accomplished by a set of superordinate programs - the emotions. They are adaptations that have arisen in response to the adaptive problem of mechanism orchestration (Tooby & Cosmides, 1990a; Tooby, 1985). In this view, the exploration of the statistical structure of ancestral situations and their relationship to the mind's battery of functionally specialized programs is central to mapping the emotions. This is because the most useful (or least harmful) deployment of programs at any given time will depend critically on the exact nature of the confronting situation.

How did emotions arise and assume their distinctive structures? Fighting, falling in love, escaping predators, confronting sexual infidelity, experiencing a failure-driven loss in status, responding to the death of a family member (and so on) each involved conditions, contingencies, situations, or event-types that recurred innumerable times in hominid evolutionary history. Repeated encounters with each kind of situation selected for adaptations that guided information-processing, behavior and the body adaptively through the clusters of conditions, demands, and contingencies that characterized that particular class of situation. This could be accomplished by engineering superordinate programs, each of which jointly mobilizes a subset of the psychological architecture's other programs in a particular configuration. Each configuration would be selected to deploy computational and physiological mechanisms in a way that, when averaged over individuals and generations, would have led to the most fitness-promoting subsequent lifetime outcome given that ancestral situation-type.

This coordinated adjustment and entrainment of mechanisms is a *mode of operation for the entire psychological architecture*, and serves as the basis for a precise computational and functional definition of each emotion state (Tooby & Cosmides, 1990a; Tooby, 1985). Each emotion entrains various other adaptive programs - deactivating some, activating others, and adjusting the modifiable parameters of still others - so that the whole system operates in a particularly harmonious and efficacious way when the individual is confronting certain kinds of triggering conditions or situations. The conditions or situations relevant to the emotions are those that (1) recurred ancestrally; (2) could not be negotiated successfully unless there was a superordinate level of program coordination (i.e., circumstances in which the independent operation of programs caused no conflicts would not have selected for an emotion program, and would lead to emotionally neutral states of mind); (3) had a rich and reliable repeated structure; (4) had recognizable cues signaling their presence; and (5) in which an error would have resulted in large fitness costs (Tooby & Cosmides, 1990a; Tooby, 1985). When a condition or situation of an evolutionarily recognizable kind is detected, a signal is sent out from the emotion program that activates the specific constellation of subprograms appropriate to solving the type of adaptive problems that were regularly embedded in that situation, and deactivates programs whose operation might interfere with solving those types of adaptive problem. Programs directed to remain active may be cued to enter subroutines that are specific to that emotion mode, and that were tailored by natural selection to solve the problems inherent in the triggering situation with special efficiency.

According to this theoretical framework, an emotion is a superordinate program whose function is to direct the activities and interactions of the subprograms governing perception; attention; inference; learning; memory; goal choice; motivational priorities; categorization and conceptual frameworks; physiological reactions (such as heart rate, endocrine function, immune function, gamete release); reflexes; behavioral decision rules; motor systems; communication processes; energy level and effort allocation; affective coloration of events and stimuli; recalibration of probability estimates, situation assessments, values, and regulatory variables (e.g., self-esteem, estimations of relative formidability, relative value of alternative goal

states, efficacy discount rate); and so on. An emotion is not reducible to any one category of effects, such as effects on physiology, behavioral inclinations, cognitive appraisals, or feeling states, because it involves evolved instructions for all of them together, as well as other mechanisms distributed throughout the human mental and physical architecture.

All cognitive programs - including superordinate programs of this kind - are sometimes mistaken for "homunculi", that is, entities endowed with "free will". A homunculus scans the environment and freely chooses successful actions in a way that is not systematic enough to be implemented by a program. It is the task of cognitive psychologists to replace theories that implicitly posit such an impossible entity with theories that can be implemented as fixed programs with open parameters. Emotion programs, for example, have a front end that is designed to detect evolutionarily reliable cues that a situation exists (whether or not these cues reliably signal the presence of that situation in the modern world); when triggered, they entrain a specific set of subprograms: those that natural selection "chose" as most useful for solving the problems that situation posed in ancestral environments. Just as a computer can have a hierarchy of programs, some of which control the activation of others, the human mind can as well. Far from being internal free agents, these programs have an unchanging structure regardless of the needs of the individual or her circumstances, because they were designed to create states that worked well in ancestral situations, regardless of their consequences in the present.

Fear (an example)

Consider the following example. The ancestrally recurrent situation is being alone at night and a situation-detector circuit perceives cues that indicate the possible presence of a human or animal predator. The emotion mode is a fear of being stalked. (In this conceptualization of emotion, there might be several distinct emotion modes that are lumped together under the folk category "fear", but that are computationally and empirically distinguishable by the different constellation of programs each entrains.) When the situation detector signals that one has entered the situation "possible stalking and ambush", the following kinds of mental programs are entrained or modified: (1) There are shifts in perception and attention: You may suddenly hear with far greater clarity sounds that bear on the hypothesis that you are being stalked, but that ordinarily you would not perceive or attend to, such as creaks or rustling. Are the creaks footsteps? Is the rustling caused by something moving stealthily through the bushes? Signal detection thresholds shift: Less evidence is required before you respond as if there were a threat, and more true positives will be perceived at the cost of a higher rate of false alarms. (2) Goals and motivational weightings change: Safety becomes a far higher priority. Other goals and the computational systems that subserve them are deactivated: You are no longer hungry; you cease to think about how to charm a potential mate; practicing a new skill no longer seems rewarding. Your planning focus narrows to the present: worries about yesterday and tomorrow temporarily vanish. Hunger, thirst, and pain are suppressed. (3) Information-gathering programs are redirected: Where is my baby? Where are others who can protect me? Is there somewhere I can go where I can see and hear what is going on better? (4) Conceptual frames shift, with the automatic imposition of categories such as "dangerous" or "safe". Walking a familiar and usually comfortable route may now be mentally tagged as "dangerous". Odd places that you normally would not occupy - a hallway closet, the branches of a tree - suddenly may become salient as instances of the category "safe" or "hiding place". (5) Memory processes are directed to new retrieval tasks: Where was that tree I climbed before? Did my adversary and his friend look at me furtively the last time I saw them? (6) Communication processes change: Depending on the circumstances, decision rules might cause you to emit an alarm cry, or be paralyzed and unable to speak. Your face may automatically assume a species-typical fear expression. (7) Specialized inference systems are activated: Information about a lion's trajectory or eye direction might be fed into systems for inferring whether the lion saw you. If the inference is yes, then a program automatically infers that the lion knows where you are; if no, then the lion does not know where you are (the "seeing-is-knowing" circuit identified by Baron-Cohen 1995, and inactive in autistics). This variable may automatically govern whether you freeze in terror or bolt. Are there cues in the lion's behavior that indicate whether it has eaten recently, and so is unlikely to be predatory in the near future? (Savanna ungulates, such as zebras and wildebeests, commonly make this kind of judgment; Marks, 1987). (8) Specialized learning systems are activated, as the large literature on fear conditioning indicates (e.g., LeDoux, 1995; Mineka & Cook, 1993; Pitman & Orr, 1995). If the threat is real, and the ambush occurs, the victim may experience an amygdala-mediated recalibration (as in post-traumatic stress disorder) that can last for the remainder of his

or her life (Pitman & Orr, 1995). (9) Physiology changes: Gastric mucosa turn white as blood leaves the digestive tract (another concomitant of motivational priorities changing from feeding to safety); adrenalin spikes; heart rate may go up or down (depending on whether the situation calls for flight or immobility), blood rushes to the periphery, and so on (Cannon, 1929; Tomaka, Blascovich, Kibler, & Ernst, 1997); instructions to the musculature (face, and elsewhere) are sent (Ekman, 1982). Indeed, the nature of the physiological response can depend in detailed ways on the nature of the threat and the best response option (Marks, 1987). (10) Behavioral decision rules are activated: Depending on the nature of the potential threat, different courses of action will be potentiated: hiding, flight, self-defense, or even tonic immobility (the latter is a common response to actual attacks, both in other animals and in humans). Some of these responses may be experienced as automatic or involuntary.

From the point of view of avoiding danger, these computational changes are crucial: They are what allowed the adaptive problem to be solved with high probability, on average over evolutionary time. Of course, in any single case they may fail, because they are only the evolutionarily-computed best bet, based on ancestrally summed outcomes; they are not a sure bet, based on an unattainable perfect knowledge of the present. Whether individuals report consciously experiencing fear is a separate question from whether their mechanisms assumed the characteristic configuration that, according to this theoretical approach, defines the fear emotion state. Individuals often behave as if they are in the grip of an emotion, while denying they are feeling that emotion. We think it is perfectly possible that individuals sometimes remain unaware of their emotion states, which is one reason we do not use subjective experience as the *sine qua non* of emotion. At present, both the function of conscious awareness, and the principles that regulate conscious access to emotion states and other mental programs are complex and unresolved questions. Mapping the design features of emotion programs can proceed independently of their resolution, at least for the present.

With the preceding view of emotions in mind, in the next section we will outline the evolutionary and cognitive principles that led us to it (detailed arguments for these positions can be found in Tooby & Cosmides, 1992, 1990a, and 1990b, and in Cosmides & Tooby 1987, 1992, 1997).

Evolutionary foundations

Chance and selection. For reasons researchers have only recently come to fully appreciate, every species has a universal, species-typical evolved architecture (Tooby & Cosmides, 1990b). These designs are largely conserved through genetic inheritance from generation to generation (accounting over the long term for homologous similarities among related species). Nevertheless, over the long run, evolutionary change takes place, and this design modification is governed by two kinds of processes: chance and selection. Random mutations are always being injected into species. What ultimately happens to each mutation is shaped both by chance and by the stable consequences the mutation has on the design of the organism -- selection. A mutational modification to a design feature will often alter how well it functions (e.g. improving the optics of the lens, or reducing a liver enzyme's detoxification efficiency). Those alterations that improve the machine's ability to solve reproduction-promoting tasks (compared to the earlier model design feature) will increase their own frequency over the generations, until (usually) they become universally incorporated into the species design. The accumulated effects of this positive feedback is one reason why species tend to have universal, species-typical evolved architecture in their functional components (see Tooby & Cosmides, 1990b for details and exceptions). Other modifications interfere with replication: these act to edit themselves from the population and the species design (negative feedback). Still others have no systematic effect: neutral alterations randomly drift in frequency, sometimes disappearing and sometimes becoming species-typical. These processes - chance and selection - explain how species acquired their designs.

For researchers seeking to understand organic design, natural selection is the most important component to consider, because it is the only force in nature that can build functional organization into organisms. Natural selection is a hill-climbing feedback process that chooses among alternative designs *on the basis of how well they function*. This is what biologists mean when they say that function determines structure. Natural selection is a causal process in which a structure spreads *because of* its functional consequences. This causal relationship is what gives theories of adaptive function their heuristic power for psychologists and biologists. If one knows what adaptive problems our ancestors faced generation after generation, one can look for mechanisms that are well-engineered for solving them.

Because of the different roles played by chance and selection, the evolutionary process builds three different types of outcomes into organisms: (1) adaptations, that is, functional machinery built by selection (usually species-typical), (2) byproducts of adaptations, which are present in the design of organisms because they are causally coupled to traits that were selected for (usually species-typical), and (3) random noise, injected by mutation and other random processes (often not species-typical) (Tooby & Cosmides, 1990a, 1990b, 1992; Williams, 1966). The emotion of sexual jealousy is an adaptation (Daly, Wilson & Weghorst, 1982; Buss, 1994); stress-induced physical deterioration is arguably a byproduct of the flight-fight system; and heritable personality variation in emotional functioning (e.g., extreme shyness, morbid jealousy, bipolar depression) is probably noise (Tooby & Cosmides, 1990b). Evidence of the presence (or absence) of high degrees of coordination between adaptive problems and the design features of putative adaptations allows researchers to distinguish adaptations, byproducts, and noise from one another (Williams, 1966; Cosmides & Tooby, 1997).

How well-designed are emotion adaptations expected to be? Organisms, as a result of millions of years of selection, are full of evolved adaptations that are improbably well-engineered to solve the adaptive problems the species encountered repeatedly during its evolution. Biologists have found that selection has routinely produced exquisitely engineered biological machines of the highest order, at all scales, from genetic error correction and quality control in protein assembly to photosynthetic pigments, the immune system, the vertebrate eye and visual system, efficient bee foraging algorithms, echolocation, and color constancy systems. While Stephen Jay Gould (1997) and his followers have energetically argued in the popular science literature that natural selection is a weak evolutionary force, evolutionary biologists, familiar with the primary literature, have found it difficult to take these arguments seriously (Tooby & Cosmides, under review).

In fact, whenever the adaptive problem can be well-specified (as in color constancy, object recognition, grammar acquisition, word meaning induction, tactile perception, chemical identification), natural computational adaptations have consistently and strikingly outperformed the best artificial devices that teams of engineers, using decades of effort, and millions of dollars of funding, have produced (consider, e.g., artificial vision or speech recognition programs). So while adaptations are in some abstract sense undoubtedly far from optimal, they are nevertheless extremely well-engineered, and their performance on the problems they evolved to solve is unrivaled by any machine yet designed by humans. The empirical evidence falsifies the claim that evolved computational adaptations tend to be crude or primitive in design, and supports the opposite view: that our mental machinery, including the emotions, are likely to be very well-designed to carry out evolved functions. For emotion researchers, this means that one's working hypotheses (which are always open to empirical revision) should begin with the expectation of high levels of evolutionary functionality, and one's research methods should be sensitive enough to detect such organization. This does not mean that emotions are well-designed for the modern world -- only that their functional logic is likely to be sophisticated and well-engineered to solve ancestral adaptive problems.

Adaptive problems. Over evolutionary time, design features are added to or discarded from the species' design because of their consequences. A design feature will cause its own spread over generations if it has the consequence of solving adaptive problems such as detecting predators, deterring sexual rivals, helping sisters, or ejecting toxin-laden food. *Adaptive problems* are evolutionarily long-enduring recurring clusters of conditions that constitute either reproductive opportunities (e.g., the arrival of a potential mate; the reflectant properties of light) or reproductive obstacles (e.g., the speed of a prey animal; the actions of a sexual rival, limited food supplies for relatives). Adaptations were designed by selection to exploit these opportunities and to circumvent these obstacles. A design feature may be said to solve an adaptive problem to the extent that its presence in an organism (when compared to alternative designs) increases the organism's net lifespan reproduction, and the reproduction of kin (who are likely to carry the same genetically based design feature; Hamilton, 1964).

Researchers less familiar with evolutionary psychology often equate adaptive problems exclusively with short-run threats to physical survival. However, survival is not central to evolution: indeed, all individual organisms die sooner or later. In contrast, genes - which can be thought of as particles of design - are potentially immortal, and design features spread by promoting the reproduction of the genes that participate in building them. Survival is significant only insofar as it promotes the reproduction of design features. It is

no more significant than anything else that promotes reproduction, and is often advantageously risked or sacrificed in the process of promoting reproduction in self, children, or other relatives.

Because events and conditions in the organism's local world are causally linked, the enhancement of its reproduction reaches out to encompass, in a network of causal linkages, all of human life, from the subtleties of facial expression to attributions of responsibility to the intrinsic rewards of projectile games to the ability to imagine alternatives. The realm of adaptive information-processing problems is not limited to one area of human life, such as sex, violence, or resource acquisition. Instead, it is a dimension cross-cutting all areas of human life, as weighted by the strange, nonintuitive metric of their cross-generational statistical effects on direct and kin reproduction.

Moreover, it is important to remember that the consequences at issue in a good design are *total lifetime fitness consequences*, not just what happens in the short run. The design features of every program have been shaped by the answer to the question: Given the situation the organism is in at each given present moment, what is the deployment at that moment of the modifiable characteristics of the individual (physiology, action, knowledge states, etc.) that will net the best return on own and kin reproduction, as accrued over the expected remainder of the individual's lifespan? Emotion programs that incline the individual to engage in seemingly pointless activities over the near-term (e.g., grief, playfulness, fascination, guilt, depression, feeling triumphant) need to be analyzed in terms of how they modify the psychological architecture for benefits that are accrued probabilistically over the long-run (e.g., gains in knowledge; recalibration of motivational priorities; the recomputation of a huge body of choice-variables in the face of information that the local world has dramatically changed).

The environment of evolutionary adaptedness. Behavior in the present is generated by evolved information-processing mechanisms that were constructed in the past because they solved adaptive problems in the ancestral environments in which the human line evolved. For this reason, evolutionary psychology is both environment-oriented and past-oriented in its functionalist orientation. Adaptations become increasingly effective as selection makes their design features more and more complementary to the long-enduring structure of the world. The articulated features of the adaptation are designed to mesh with the features of the environment that were stable during the adaptation's evolution, so that their interaction produced functional outcomes. The regulation of breathing assumes the presence of certain long-enduring properties of the atmosphere and the respiratory system. Vision assumes the presence of certain evolutionarily stable properties of surfaces, objects, and terrestrial spectral distributions. The lactase digestive enzyme presupposes an infant diet of milk with lactose. And each emotion presupposes that certain cues signal the presence of a structure of events and conditions that held true during the evolution of that emotion. Disgust circuits presume a world in which rotten smells signal toxins or microbial contamination, for example.

Accordingly, to understand an adaptation as a problem-solver, one needs to model the enduring properties of the task-environment that selected for that adaptation -- the *environment of evolutionary adaptedness*, or EEA. Although the hominid line is thought to have first differentiated from the chimpanzee lineage on the African savannas and woodlands, the EEA is not a place or time. It is the statistical composite of selection pressures that caused the genes underlying the design of an adaptation to increase in frequency until they became species-typical or stably persistent (Tooby & Cosmides, 1990a). Thus, statistical regularities define the EEA for any given adaptation. The conditions that characterize the EEA are usefully decomposed into a constellation of specific environmental regularities that had a systematic (though not necessarily unvarying) impact on reproduction and that endured long enough to work evolutionary change on the design of an adaptation. These regularities can include complex conditionals (e.g., if one is male hunter-gatherer *and* one is having a sexual liaison with someone else's mate *and* that is discovered, then one is the target of lethal retributory violence 37% of the time). Descriptions of these statistical regularities are essential parts of the construction of a task analysis of the adaptive problem a hypothesized adaptation evolved to solve (Tooby & Cosmides, 1990a). Conceptualizing the EEA in statistical terms is fundamental to the functional definition of emotion that we presented above and will elucidate below.

Cognitive foundations

The cognitive science resolution of the mind-body problem. Evolutionary psychology starts with a fundamental insight from cognitive psychology: The brain is a machine designed to process information. From this perspective, one can define the *mind* as a set of information-processing procedures (cognitive programs) that are physically embodied in the neural circuitry of the brain. For cognitive scientists, *brain* and *mind* are terms that refer to the same system, which can be described in two complementary ways - either in terms of its physical properties (the neural), or in terms of its information-processing operation (the mental). The mind is what the brain does, described in computational terms (Jackendoff, 1987; Cosmides & Tooby, 1987; Pinker, 1997). This approach allows mental operations to be described with great precision: one is led to specify what information is extracted from the environment, what procedures act to transform it, what formats are involved in its representation or storage, and what operations access it to govern decision-making, physiological or behavioral regulation, or further information integration. Also, because it provides an intelligible way of relating physical and mental phenomena, discoveries in brain science (e.g., from dissociation studies and neuroimaging) can be used in making inferences about the mind, and vice versa - a process that is leading to a principled mapping between brain and mind (for reviews, see, Gazzaniga, 1995, 2000).

An evolutionary perspective makes clear why the cognitive or computational level of description is more than an analogy. Whereas other parts of the body were designed for lifting loads, grinding food, chemically extracting nutrients, and so on, the brain was designed by evolution to use information derived from the environment and the body to functionally regulate behavior and the body. The brain came into existence and, over evolutionary time, accreted its present complex structure because, in ancestral populations, mutations that created or altered cognitive programs such that they more successfully carried out adaptively consequential information-processing tasks were differentially retained, replicated, and incorporated into our species' neural design.

The ancestral world posed recurrent information-processing problems, such as What substances are best to eat? or What is the relationship between others' facial expressions and their mental states? Information-processing programs - food preferences and aversions, or rules for inferring emotions from facial expressions - acquired one set of design features rather than many others because the retained features better computed solutions to these information-processing problems. Over evolutionary time, it was the *computational* properties of alternative neural circuits - their relative ability to solve adaptive information-processing tasks - that caused some neural circuits to be selected for, and others to be selected out. So, from an evolutionary and a functional point of view, the brain is intrinsically and by its nature an organ of computation - a set of information-processing devices realized in neural tissue (Cosmides & Tooby, 1987; Tooby & Cosmides, 1992; Pinker, 1997). A key task for psychologists, then, is to discover, inventory and map the "circuit logic" of the collection of programs that constitute the human mind, and relate how that adaptive logic maps onto the suite of informational problems faced by our hunter-gatherer ancestors.

Emotion and computation. It may strike some as odd to speak about love or jealousy or disgust in computational terms. "Cognition" and "computation" have affectless, flavorless connotations. In everyday language, the term "cognition" is often used to refer to a particular *subset* of information-processing - roughly, the effortful, conscious, voluntary, deliberative kind of thinking one does when solving a mathematics problem or playing chess: what is sometimes called "cold cognition". This use of "cognition" falls out of the folk psychological classification of thinking as distinct from feeling and emotion, and it appears in a few subfields of psychology as well (particularly those concerned with education and the acquisition of skills that must be explicitly taught). As a result, one sometimes sees articles in the psychological literature on how emotion, affect, or mood influence "cognition".

However, from an evolutionary cognitive perspective, one cannot sensibly talk about emotion affecting cognition because cognition refers to a language for describing all of the brain's operations, including emotions and reasoning (whether deliberative or nonconscious), and not to any particular subset of operations. If the brain evolved as a system of information-processing relations, then emotions are, in an evolutionary sense, best understood as information-processing relations - i.e., programs - with naturally selected functions. Initially, the commitment to exploring the underlying computational architecture of the emotions may strike one as odd or infelicitous, but it leads to a large number of scientific payoffs, as we will sketch out below.

Thus, the claim that emotion is computational does not mean that an evolutionary psychological approach reduces the human experience to bloodless, affectless, disembodied ratiocination. Every mechanism in the brain - whether it does something categorizable as "cold cognition" (such as inducing a rule of grammar or judging a probability) or as "hot cognition" (such as computing the intensity of parental fear, the imperative to strike an adversary, or an escalation in infatuation) - depends on an underlying computational organization to give its operation its patterned structure, as well as a set of neural circuits to physically implement it.

Of course, shifting terminology (e.g., from "cognition" as thinking to "cognition" as everything mental) does nothing to invalidate research done with the old terminology, and valuable research has been done exploring how various emotion states modify performance on tasks that require deliberative thinking (e.g., Isen, 1987; Mackie & Worth, 1991). But an evolutionary and computational view of emotion can open up for exploration new empirical possibilities obscured by other frameworks. An evolutionary perspective breaks categories like "thinking" into a large set of independent domain-specific programs, and so opens up the possibility that distinct emotions affect separate inference programs in diverse, yet functionally patterned, ways, rather than in a single, aggregate way.

Domain specificity and functional specialization. A basic engineering principle is that the same machine is rarely capable of solving two different problems equally well. Cork-screws and cups have different properties because they are solutions to different problems, and each therefore solves its targeted problem better than the other. Similarly, natural selection has constructed different tissues and organs, like the heart for pumping blood and the liver for detoxifying poisons, for exactly this reason. This same principle applies to our evolved cognitive programs and neural circuitry. Different information-processing problems usually require different procedures for their successful solution. For example, to solve the adaptive problem of selecting a good mate, one's choices must be guided by qualitatively different standards than when choosing the right food, or the right habitat, or the right meaning for an unfamiliar word. Implementing different solutions requires different, functionally distinct mechanisms (Sherry & Schacter, 1987; Gallistel, 1995). Speed, reliability, and efficiency can be engineered into specialized mechanisms because they do not need to make trade-offs between mutually incompatible task demands, and because they can use problem-solving principles that work in one domain but not in others. (For detailed arguments, both on the weakness of domain-general architectures and on the many advantages of architectures that include a large number of domain-specific computational devices, see Cosmides & Tooby, 1987, 1994; Tooby & Cosmides, 1990a, 1992).

The application of these principles to the design of the mind has convinced many scientists, including most evolutionary psychologists, that the human cognitive architecture is *multimodular*: that it is composed of a large number of information-processing programs, many of which are functionally specialized for solving a different adaptive problem. These adaptations appear to be domain-specific expert systems, equipped with "crib sheets": inference procedures, regulatory rules, motivational priorities, goal-definitions, and assumptions that embody knowledge, regulatory structure, and value weightings specific to an evolved problem domain. These generate correct (or, at least, adaptive) outputs that would not be warranted on the basis of perceptual data processed through some general-purpose decisional algorithm. In the last two decades, many cognitive researchers have found evidence for the existence a diverse collection of inference systems, including specializations for reasoning about objects, physical causality, number, language, the biological world, the beliefs and motivations of other individuals, and social interactions (for reviews, see Hirschfeld & Gelman, 1994; *Cognitive Science*, volume 14 (1990), and Barkow, Cosmides and Tooby, 1992). These domain-specific inference systems have a distinct advantage over domain-independent ones, akin to the difference between experts and novices: experts can solve problems faster and more efficiently than novices because they already know a lot about the problem domain, and because they are equipped with specialized tools and practices.

Each adaptive problem recurred millions of times in the EEA, and so manifested a statistical and causal structure whose elements were available for specialized exploitation by design features of the evolving adaptation. For example, predators use darkness and cover to ambush. Physical appearance varies with fertility and health. Children one's mother regularly feeds are usually genetic siblings. Specialized programs - for predator fear, sexual attraction, and incest avoidance, respectively - could evolve whose configuration

of design features embodied and/or exploited these statistical regularities, allowing these adaptive problems to be solved economically, reliably, and effectively. Such specializations, by embodying "innate knowledge" about the problem space, operate better than any general learning strategy could. A child did not have to wait to be ambushed and killed in the dark to prudently modulate her activities. Adults did not need to observe the negative effects of incest, because the Westermarck mechanism mobilizes disgust towards having sex with likely siblings (Shepher 1983).

Selection detects the individually unobservable. Animals subsist on information. The single most limiting resource to reproduction is not food or safety or access to mates, but what makes them each possible: the information required for making adaptive behavioral choices. However, many important features of the world cannot be perceived directly. Cognitive adaptations can use perceivable events as cues for inferring the status of important, nonperceivable sets of conditions, provided there was a predictable probabilistic relationship between them that was maintained over evolutionary time. Natural selection can extract statistical relationships that would be undetectable to any individual organism (Cosmides & Tooby, 1987; Tooby & Cosmides, 1990a). It does this by testing randomly generated alternative designs, each of which embodies different assumptions about the structure of the world, and retaining the ones that succeed most effectively. The most effective design will be the one that best embodies design features that reflect most closely the actual long-term statistical structure of the ancestral world. Designs whose features exploited these real but ontogenetically unobservable relationships outperformed those that depended on different relationships, or that only responded to conditions the individual could observe during her lifetime.

This is why tabula rasa models of human and nonhuman minds are evolutionarily impossibilities (Cosmides & Tooby, 1987). For example, the negative effects of incestuous conceptions are difficult for any individual to observe in the absence of a modern controlled study with numerous participants, much less integrate rationally into one's motivational system. Fortunately, the consequences of incest over evolutionary time selected for specialized disgust mechanisms that reflected the ancestral distribution of choice-consequence pairings, and so are designed to guide humans away from incestuous unions between fertile adults, given appropriate cues of familial connection such as co-residence in the first years of life (Shepher, 1983). Evolved psychological adaptations are selected to use cues that (1) can be reliably and easily detected by the individual, and (2) reliably predicted the hidden structure of conditions relevant to determining which course of action one should take.

The functional structure of an emotion program evolved to match the evolutionarily summed structure of its target situation. The set of human emotion programs assumed their evolved designs through interacting with the statistically defined structure of human environments of evolutionary adaptedness. Each emotion program was constructed by a selective regime imposed by a particular evolutionarily recurrent situation. By an evolutionarily recurrent situation, we mean a cluster of repeated probabilistic relationships among events, conditions, actions, and choice-consequences, that endured over a sufficient stretch of evolutionary time to have had selective consequences on the design of the mind, and that were probabilistically associated with cues detectable by humans.

For example, the condition of having a mate plus the condition of one's mate copulating with someone else constitutes a situation of sexual infidelity: a situation that has recurred over evolutionary time, even though it has not happened to every individual. Associated with this situation were cues reliable enough to allow the evolution of a "situation detector" (e.g., observing a sexual act, flirtation, or even the repeated simultaneous absence of the suspected lovers are cues that could trigger the categorization of a situation as one of infidelity). Even more importantly, there were many necessarily or probabilistically associated elements that tended to be present in the situation of infidelity as encountered among our hunter-gatherer ancestors. Additional elements include: (1) a sexual rival with a capacity for social action and violence, as well as allies of the rival; (2) a discrete probability that one's mate has conceived with the sexual rival; (3) changes in the net lifetime reproductive returns of investing further in the mating relationship; (4) a probable decrease in the degree to which the unfaithful mate's mechanisms value the victim of infidelity (the presence of an alternative mate lowers replacement costs); (5) a cue that the victim of the infidelity will likely have been deceived about a range of past events, leading the victim to confront the likelihood that his or her memory is permeated with false information; (6) the victim's status and reputation for being effective at defending his or her interests in general would be likely to plummet, inviting challenges in other arenas.

These are just a few of the many factors that constitute a list of elements associated in a probabilistic cluster, and that constitute the evolutionary recurrent structure of a situation of sexual infidelity. The emotion of sexual jealousy evolved in response to these properties of the world, and there should be evidence of this in its computational design.

Emotion programs have evolved to take such elements into account, whether they can be perceived or not. Thus, not only do cues of a situation trigger an emotion mode, but embedded in that emotion mode is a way of seeing the world and feeling about the world related to the ancestral cluster of associated elements. Depending on the intensity of the jealousy evoked, less and less evidence will be required for an individual to believe that these conditions apply to their personal situation. Individuals with morbid jealousy, for example, may hallucinate counterfactual but evolutionarily thematic contents.

To the extent that situations exhibit a structure repeated over evolutionary time, their statistical properties will be used as the basis for natural selection to build an emotion program whose detailed design features are tailored for that situation. This is accomplished by selection, acting over evolutionary time, differentially incorporating program components that dovetail with individual items on the list of properties probabilistically associated with the situation.

For example, if in ancestral situations of sexual infidelity there was a substantially higher probability of a violent encounter than in its absence, then the sexual jealousy program will have been shaped by the distillation of those encounters, and the jealousy subroutines will have been adjusted to prepare for violence in proportion to the raised probability in the ancestral world. (Natural selection acts too slowly to have updated the mind to post-hunter-gatherer conditions.) Each of these sub-elements and the adaptive circuits they require can be added together to form a general theory of sexual jealousy.

The emotion of sexual jealousy constitutes an organized mode of operation specifically designed to deploy the programs governing each psychological mechanism so that each is poised to deal with the exposed infidelity: physiological processes are prepared for such things as violence, sperm competition, and the withdrawal of investment; the goal of deterring, injuring, or murdering the rival emerges; the goal of punishing, deterring, or deserting the mate appears; the desire to make oneself more competitively attractive to alternative mates emerges; memory is activated to reanalyze the past; confident assessments of the past are transformed into doubts; the general estimate of the reliability and trustworthiness of the opposite sex (or indeed everyone) may decline; associated shame programs may be triggered to search for situations in which the individual can publicly demonstrate acts of violence or punishment that work to counteract an (imagined or real) social perception of weakness; and so on.

It is the relationship between the summed details of the ancestral condition and the detailed structure of the resulting emotion program that makes this approach so useful for emotion researchers. Each functionally distinct emotion state - fear of predators, guilt, sexual jealousy, rage, grief, and so on - will correspond to an integrated mode of operation that functions as a solution designed to take advantage of the particular structure of the recurrent situation or triggering condition to which that emotion corresponds. This approach can be used to create theories of each individual emotion, through three steps: (1) Reconstructing the clusters of properties of ancestral situations; (2) Constructing engineering analyses about how each of the known or suspected psychological mechanisms in the human mental architecture should be designed to deal with each ancestral condition or cluster of conditions, and integrating these into a model of the emotion program; (3) Constructing or conducting experiments and other investigations to test, and revise the models of emotion programs.

It is also important to understand that evolutionarily recurrent situations can be arrayed along a spectrum in terms of how rich or skeletal is the set of probabilistically associated elements that defines the situation. A richly structured situation, such as sexual infidelity or predator ambush will support a richly substructured emotion program in response to the many ancestrally correlated features: Many detailed adjustments will be made to many psychological mechanisms as instructions for the mode of operation. In contrast, some recurrent situations have less structure (i.e., they share fewer properties in common), and so the emotion mode makes fewer highly specialized adjustments, imposes fewer specialized and compelling interpretations and behavioral inclinations, and so on. For example, surges of happiness or joy are an emotion program that

evolved to respond to the recurrent situation of encountering unexpected positive events (as will be explained). The class of events captured by "unexpectedly positive" is extremely broad, general, and have only a few additional properties in common. Emotion programs at the most general and skeletal end of this spectrum correspond to what some call "mood" (happiness, sadness, excitement, anxiety, playfulness, homesickness, and so on).

How to characterize an emotion

To characterize an emotion adaptation, one must identify the following properties of environments and of mechanisms.

1) *An evolutionarily recurrent situation or condition*: A repeated structure of environmental and organismic properties, characterized as a complex statistical composite of how such properties covaried in the environment of evolutionary adaptedness. Examples of situations are being in a depleted nutritional state, competing for maternal attention, being chased by a predator, being about to ambush an enemy, having few friends, experiencing the death of a spouse, being sick, having experienced a public success, having others act in a way that damages you without regard for your welfare, having injured a valued other through insufficient consideration or self-other behavioral trade-offs; having a baby.

2) *The adaptive problem*: The identification of which organismic states and behavioral sequences will lead to the best average functional outcome for the remainder of the lifespan, given the situation or condition. For example, what is the best course of action when others take the products of your labor without your consent? What is the best course of action when you are in a depleted nutritional state? What is the best course of action when a sibling makes a sexual approach?

3) *Cues that signal the presence of the situation*: For example, low blood sugar signals a depleted nutritional state, the looming approach of a large fanged animal signals the presence of a predator, seeing your mate having sex with another signals sexual infidelity; finding yourself often alone, rarely the recipient of beneficent acts, or actively avoided by others signals that you have few friends.

4) *Situation-detecting algorithms*. A multimodular mind must be full of "demons": algorithms that detect situations. *The New Hacker's Dictionary* defines a "demon" as a "portion of a program that is not invoked explicitly, but that lies dormant waiting for some condition(s) to occur." (Raymond, 1991, p. 124). Situation-detecting subprograms lie dormant, until they are activated by a specific constellation of cues that precipitates the analysis of whether a particular ancestral situation has arisen. If the assessment is positive, it sends the signal that activates the associated emotion program. Emotion demons need two kinds of subroutines:

(a) *Algorithms that monitor for situation-defining cues*: These include perceptual mechanisms, proprioceptive mechanisms, and situation-modeling memory. They take the cues in (3) as input.

(b) *Algorithms that detect situations*: these programs take the output of the monitoring algorithms and targeted memory registers in (a) as input, and through integration, probabilistic weighting, and other decision criteria, identify situations as absent or present with some probability.

The assignment of a situation interpretation to present circumstances involves a problem in signal detection theory (Swets, Tanner & Birdsall 1964; see also Gigerenzer & Murray 1987). Animals should be designed to "detect" what situation they are in on the basis of cues, stored variables, and specialized interpretation algorithms. Selection will not shape decision rules so that they act solely on the basis of what is most likely to be true, but rather on the basis of the weighted consequences of acts given that something is held to be true. Should you walk under a tree that might conceal a predator? Even if the algorithms assign a 51% (or even 98%) probability to the tree being predator-free, under most circumstances an evolutionarily well-engineered decision rule should cause you to avoid the tree - to act as if the predator were in it. The benefits of calories saved via a shortcut, scaled by the probability that there is no predator in the tree, must be weighed against the benefits of avoiding becoming catfood, scaled by the probability that there is a predator in the tree. Because the costs and benefits of false alarms, misses, hits, and correct rejections are often unequal, the decision rules may still treat as true situations that are unlikely to be true. In the modern world,

this behavior may look "irrational" (as is the case with many phobias), but we do it because such decisions were adaptive under ancestral conditions.

Situation detecting algorithms can be of any degree of complexity, from demons that monitor single cues (e.g., "snake present") to algorithms that carry out more complex cognitive assessments of situations and conditions (LeDoux, 1995; Lazarus & Lazarus, 1994; Tooby & Cosmides, 1990a). Inherent in this approach is the expectation that the human mind has a series of evolved subsystems designed to represent events in terms of evolutionarily recurrent situations and situational subcomponents. The operation of these representational systems are not necessarily consciously accessible. By their structure, they impose an evolutionary organization on representational spaces that are updated by data inputs. When the representational space assumes certain configurations, an interpretation is triggered that activates the associated emotion program - corresponding approximately to what others have called a cognitive appraisal (see, e.g., Lazarus & Lazarus, 1994). It is important to recognize that the evolutionary past frames the experienced present, because these situation-detecting algorithms provide the dimensions and core elements out of which many cross-culturally recurring representations of the world are built. To some extent, the world we inhabit is shaped by the continuous interpretive background commentary provided by these mechanisms.

5) *Algorithms that assign priorities*: A given world-state may correspond to more than one situation at a time, for example, you may be nutritionally depleted *and* in the presence of a predator. The prioritizing algorithms define which emotion modes are compatible (e.g., hunger and boredom), and which are mutually exclusive (e.g., feeding and predator escape). Depending on the relative importance of the situations and the reliability of the cues, the prioritizing algorithms decide which emotion modes to activate and deactivate, and to what degree. Selection, through ancestral mutant experiments, would have sorted emotions based on the average importance of the consequences stemming from each, and the extent to which joint activation was mutually incompatible (or facilitating). (Prioritizing algorithms can be thought of as a supervisory system operating over all of the emotions.)

6) *An internal communication system*: Given that a situation has been detected, the internal communication system sends a situation-specific signal to all relevant programs and mechanisms; the signal switches them into the appropriate adaptive emotion mode. In addition, information is fed back into the emotion program from other programs and systems that assess body states, which may govern the intensity, trajectory, supplantation, or termination of the emotion.

Some modes of activation of the cognitive system are accompanied by a characteristic feeling state, a certain quality of experience. The fact that we are capable of becoming aware of certain physiological states - our hearts thumping, bowels evacuating, stomach tightening - is surely responsible for some of the qualia evoked by emotion states that entrain such responses. The fact that we are capable of becoming aware of certain mental states - such as retrieved memories of past events - is probably responsible for other qualia. But it is also possible that in some cases, the characteristic feeling state that accompanies an emotion mode results (in part) from mechanisms that allow us to sense the *signal* that activates and deactivates the relevant programs. Such internal sensory mechanisms - a kind of cognitive proprioception - can be selected for if there are mechanisms that require as input the information that a particular emotion mode has been activated. (This might be true, for example, of mechanisms designed to inhibit certain stimulus driven actions when the conditions are not auspicious.)

7) *Each program and physiological mechanism entrained by an emotion program must have associated algorithms that regulate how it responds to each emotion signal*: These algorithms determine whether the mechanism should switch on or switch off, and if on, what emotion-specialized performance it will implement. For example, there should be algorithms in the auditory system that, upon detecting the fear signal (see (6)), reset signal detection thresholds, increasing acuity for predator-relevant sounds.

What kinds of programs can emotions mobilize?

Any controllable biological process that, by shifting its performance in a specifiable way, would lead to enhanced average fitness outcomes should have come to be partially governed by emotional state (see (7))

above). Such processes include:

Goals: The cognitive mechanisms that define goal-states and choose among goals in a planning process should be influenced by emotions. For example, vindictiveness—a specialized subcategory of anger—may define "injuring the offending party" as a goal state to be achieved. (Although the evolved functional logic of this process is deterrence, this function need not be represented, either consciously or unconsciously, by the mechanisms that generate the vindictive behavior.)

Motivational priorities: Mechanisms involved in hierarchically ranking goals or calibrating other kinds of motivational and reward systems should be emotion-dependent. What may be extremely unpleasant in one state, such as harming another, may seem satisfying in another state (e.g., aggressive competition may facilitate counter-empathy). Different evolutionarily recurrent situations predict the presence—visible or invisible—of different opportunities, risks, and payoffs, so motivational thresholds and valences should be entrained. For example, a loss of face should increase the motivation to take advantage of opportunities for status-advancement, and decrease attention to attendant costs.

Information-gathering motivations: Because establishing which situation you are in has enormous consequences for the appropriateness of behavior, the process of detection should in fact involve specialized inference procedures and specialized motivations to discover whether certain suspected facts are true or false. What one is curious about, what one finds interesting, what one is obsessed with discovering should all be emotion-specific.

Imposed conceptual frameworks: Emotions should prompt construals of the world in terms of concepts that are appropriate to the decisions that must be made. When angry, domain-specific concepts such as social agency, fault, responsibility and punishment will be assigned to elements in the situation. When hungry, the food-nonfood distinction will seem salient. When endangered, safety-categorization frames will appear. The world will be carved up into categories based partly on what emotional state an individual is in.

Perceptual mechanisms: Perceptual systems may enter emotion-specific modes of operation. When fearful, acuity of hearing may increase. Specialized perceptual inference systems may be mobilized as well: If you've heard rustling in the bushes at night, human and predator figure-detection may be particularly boosted, and not simply visual acuity in general. In fact, non-threat interpretations may be depressed, and the same set of shadows will "look threatening"—that is, given a specific threatening interpretation such as "a man with a knife"—or not, depending on emotion-state.

Memory: The ability to call up particularly appropriate kinds of information out of long term memory ought to be influenced. A woman who has just found strong evidence that her husband has been unfaithful may find herself flooded by a torrent of memories about small details that seemed meaningless at the time but that now fit into an interpretation of covert activity. We also expect that what is stored about present experience will also be differentially regulated. Important or shocking events, for example, may be stored in great detail (as has been claimed about "flashbulb memories", for example), but other more moderate emotion-specific effects may occur as well.

Attention: The entire structure of attention, from perceptual systems to the contents of high level reasoning processes, should be regulated by emotional state. If you are worried that your spouse is late and might have been injured, it is hard to concentrate on other ongoing tasks (Derryberry & Tucker, 1994), but easy to concentrate on danger-scenarios. Positive emotions may broaden attentional focus (Fredrickson, 1998).

Physiology: Each organ system, tissue, or process is a potential candidate for emotion-specific regulation, and "arousal" is insufficiently specific to capture the detailed coordination involved. Each emotion program should send out a different pattern of instructions (to the face and limb muscles, the autonomic system, etc.) to the extent that the problems embedded in the associated situations differ. This leads to an expectation of that different constellations of effects will be diagnostic of different emotion states (Ekman, Levenson & Friesen 1983). Changes in circulatory, respiratory, and gastrointestinal functioning are well-known and documented, as are changes in endocrinological function. We expect thresholds regulating the contraction of various muscle groups to change with certain emotion states, reflecting the probability that they will need to

be employed. Similarly, immune allocation and targeting may vary with disgust, with the potential for injury, or with the demands of extreme physical exertion.

Communication and emotional expressions: Emotion programs are expected to mobilize many emotion-specific effects on the subcomponents of the human psychological architecture relevant to communication. Most notably, many emotion programs produce characteristic species-typical displays that broadcast to others the emotion state of the individual (Ekman, 1982). Ekman and his colleagues have established in a careful series of landmark studies that many emotional expressions are human universals, both generated and recognized reliably by humans everywhere they have been tested (Ekman, 1994). Indeed, many emotional expressions appear to be designed to be informative, and these have been so reliably informative that humans co-evolved automated interpreters of facial displays of emotion that decode these public displays into knowledge of others' mental states. It is surely true that people sometimes "lie" with their faces. But programs for inferring emotion states from facial displays would not have evolved unless doing so created a net advantage for the inferer, suggesting that these inferences were warranted more often than not.

Two things are communicated by an authentic emotional expression: (1) that the associated emotion program has been activated in an individual, providing observers with information about the state of that individual's mental programs and physiology (e.g., I am afraid); and (2) the identity of the evolutionarily recurrent situation being faced, in the estimation of the signaler (e.g., the local world holds a danger). Both are highly informative, and emotional expressions provide a continuous commentary on the underlying meaning of things to companions. This provokes the question: Why did selection build facial, vocal and postural expressions at all? More puzzlingly, why are they often experienced as automatic and involuntary?

From an evolutionary perspective, sometimes it is beneficial to provide information to others, and other times it is injurious, so most evolved communication systems involve close regulation of whether to transmit information or not. Usually this leads to a system, like language, in which the decision to communicate something (or not) can be made by the individual in detailed response to the immediate circumstances. The apparent selective disadvantages of honestly and automatically broadcasting one's emotional state have led Fridlund (1994), for example, to argue that expressions must be voluntary and intentional communications largely unconnected to emotion state. Undoubtedly they sometimes are. But even when people deliberately lie, microexpressions of face and voice often leak out (Ekman, 1985), suggesting that certain emotion programs do in fact create involuntarily emitted signals that reliably broadcast the person's emotion state. Why?

Natural selection has shaped emotion programs to signal their activation, or not, on an emotion by emotion basis. Summing for each emotion program considered by itself (jealousy, loneliness, disgust, predatoriness, parental love, sexual attraction, gratitude, fear), there was a net benefit or cost to having others know that mental state, averaged across individuals over evolutionary time. For those recurrent situations in which, on average, it was beneficial to share one's emotion state (and hence assessment of the situation) with those one was with, species-typical facial and other expressions of emotion were constructed by selection. For example, fear was plausibly beneficial to signal, because it signaled the presence of a danger that might also menace one's kin and cooperators, and also informed others in a way that might recruit assistance.

Nevertheless, averaged over evolutionary time, it was functional for the organism to signal the activation of only *some* emotion states. The conditions favoring signaling an emotion are hard to meet, so only some emotions out of the total species-typical set are associated with distinctive, species-typical facial expressions. There should be a larger set of emotions that have no automatic display. Jealousy, guilt, or boredom are all genuine emotions lacking distinctive signals. This changes the question from: Why are emotions automatically signaled? to Why are *some* emotions automatically signaled? When selection is neutral, the signs of an emotion should only be the byproducts of whatever is necessary to run the emotion program, without any selection to make the cues informative. When selection disfavors others knowing the organism's internal state, selection should suppress and obscure external cues identifying internal states. Precisely because they publicly signal themselves, our attention goes disproportionately to the subset of emotions that do come equipped with emotional expressions. We think it likely that this had an impact on the history of emotion research.

Three factors govern whether transmitting information will be beneficial or harmful: the signaler's relationship to the audience, the nature of the information that an emotion signal would release, and the computational overhead of computing the benefits and costs of information sharing on a case by case basis in order to regulate whether to make a broadcast (Tooby & Cosmides, 1996). In general (but with some notable exceptions), the closer the cooperative relationship and shared fitness interests, the more beneficial it is to share information; the more distant and adversarial the relationship, the more harmful it is. For this reason, we expect that circuits have evolved that regulate global emotional expressiveness depending on whether one is (apparently) alone, with people one shares interests with, or with social antagonists such as enemies or higher ranking individuals where leakage of damaging information should be suppressed. This global regulation may be largely automatic and nonconscious, and may involve open parameters set culturally and developmentally. Other things being equal, individuals will be shyer and less spontaneous with strangers (creating problems in public speaking), and more expressive with intimates. Similarly, it may be that male-female differences in emotional expressiveness arise from an evolutionary history in which males were on average more often in the presence of potential adversaries. Of course, it is beneficial to the transmitter to share certain types of information with adversaries, such as anger, triumph, or surrender, but many other types (fear of adversaries, pain, anxiety about weaknesses) ought to be suppressed.

The nature of the information broadcast has two components: (1) reliable consequences, predicted by the identity of the emotion, and (2) context-specific consequences (Tooby & Cosmides, 1996). The first component can be handled by automating the broadcast of the identity of those emotions that, on average, reliably produced a benefit when shared: Approval or disapproval assist in communicating to social interactants one's values; fear communicated the nature of a common danger; disgust communicated avoidance and spoilage; anger, a conflict of values with a willingness to enforce one's values with a sanction. The second, context-specific component requires computational circuitry to calculate the consequences of releasing a piece of information into the social world - a very complex set of computations. The benefit gained by inhibiting release of an expression on a case by case basis must be large enough to offset the cost of such computations for selection to favor the evolution of such regulatory circuits. The overall result of these selection pressures would be that some emotions would evolve to be automatically broadcast, others would not evolve a signal, and a third category would evolve circuits that regulate the broadcast to some extent, just as it is in language.

Nevertheless, the automatic, involuntary expression of many emotions is a key feature of the biology and social life of our species, and their presence provides powerful evidence that ancestral humans spent a large portion of their time with close cooperators, as opposed to antagonists and competitors. Indeed, species ought to vary in the magnitude of automatic emotion signaling, and in which emotions are signaled, based on the social ecology of the species. Highly cooperative social species, such as canids, are expected to (and appear to) have a rich repertoire of emotion signals, while more solitary species, such as felids, should have fewer emotion signals.

Behavior: All psychological mechanisms are involved in the generation and regulation of behavior, so obviously behavior will be regulated by emotion state. More specifically, however, mechanisms proximately involved in the generation of actions (as opposed to processes like face recognition that are only distally regulatory) should be very sensitive to emotion state. Not only may highly stereotyped behaviors of certain kinds be released (as during sexual arousal or rage, or as with species-typical facial expressions and body language), but more complex action-generation mechanisms should be regulated as well. Specific acts and courses of action will be more available as responses in some states than in others, and more likely to be implemented. Emotion mode should govern the construction of organized behavioral sequences that solve adaptive problems.

Biologists, psychologists, and economists who adopt an evolutionary perspective have recognized that many forms of social interactions can be modeled using game theory (Maynard Smith 1982). If the EEA imposes certain evolutionarily repeated games, then the "strategies" (the evolved cognitive programs that govern behavior in those contexts) should evolve in the direction of choices that lead to the best expected fitness payoffs. The strategy activated in the individual should match the game (e.g., exchange) and the state of play in the game (e.g., just been cheated), a process that requires the system of cues, situation-detection, and so on, already discussed. So, different emotion and inference programs or subprograms may have evolved to

correspond to various evolved games, including zero sum competitive games, positive sum exchange games, coalitional lottery games, games of aggressive competition corresponding to "chicken", and so on (for exchange, see Cosmides, 1989; Cosmides & Tooby, 1992). Corresponding emotion programs guide the individual into the appropriate interactive strategy for the social "game" being played, given the state of play. Surprisingly, for some games, rigid obligatory adherence to a prior strategy throughout the game is better than the ability to revise and change strategies ("voluntarily") in the light of events. If an individual contemplating a course of action detrimental to you knew you would take revenge, regardless of the magnitude of the punishment to you that they might unleash, then they are less likely to take such harmful action. This may translate into emotion programs in which the desire to attempt certain actions should be overwhelming, to the point where the actions are experienced as compulsory. In the grip of such programs, competing programs, including the normal integration of prudential concerns and social consequences, are muted or terminated. For example, the desire to avenge a murder or an infidelity is often experienced in this way, and are even culturally recognized as "crimes of passion" (Daly & Wilson, 1988). In modern state societies, where there are police who are paid to punish and otherwise enforce agreements, it is easy to underestimate the importance that deterrence based on the actions of oneself and one's coalition had in the Pleistocene. Hirshleifer (1987) and Frank (1988) are evolutionary economists who have pursued this logic the furthest, arguing that many social behaviors are the result of such "commitment problems".

Specialized inference: Research in evolutionary psychology has shown "thinking" and reasoning is not a unitary category, but is carried out by a variety of specialized mechanisms. So, instead of emotion activating or depressing "thinking" in general, the specific emotion program activated should *selectively* activate appropriate specialized inferential systems, such as cheater detection (Cosmides 1989; Cosmides & Tooby 1989, 1992), bluff detection (Tooby & Cosmides, 1989), precaution detection (Fiddick, Cosmides & Tooby, in press), attributions of blame and responsibility, and so on. We are presently conducting research to see whether, as predicted, fear influences precautionary reasoning, competitive loss regulates bluff detection, and so on.

Reflexes: Muscular coordination, tendency to blink, threshold for vomiting, shaking, and many other reflexes are expected to be regulated by emotion programs to reflect the demands of the evolved situation.

Learning: Emotion mode is expected to regulate learning mechanisms. What someone learns from stimuli will be greatly altered by emotion mode, because of attentional allocation, motivation, situation-specific inferential algorithms, and a host of other factors. Emotion mode will cause the present context to be divided up into situation-specific functionally appropriate categories so that the same stimuli and the same environment may be interpreted in radically different ways, depending on emotion state. For example, which stimuli are considered similar should be different in different emotion states, distorting the shape of the individual's psychological "similarity space" (Shepard 1987). Highly specialized learning mechanisms might be activated, such as those that control food aversions (Garcia, 1990) or predator learning (Mineka & Cooke, 1985), or fear conditioning (LeDoux, 1995). Happiness is expected to signal the energetic opportunity for play, and allow other exploratory agendas to be expressed (Frederickson, 1998).

Affective coloration of events and stimuli as a form of learning: A behavioral sequence is composed of many acts. Each of these acts can be thought of as an intermediate "factor" in the production of a behavioral sequence (to use economic terminology). Determining which courses of action are worthwhile and which are not is a major informational problem. The payoff of each "factor of production"-of each act in the sequence-must be computed before an agent could determine whether the whole sequence was worthwhile. Every time there is a change in the world (e.g., death of a spouse, the acquisition of a better foraging tool) that affects the probable payoff of an act, or new information that allows a better evaluation of payoffs, this value needs to be recomputed. Evaluating entire chains as units is not sufficient, because each item in a chain (staying behind from the hunt, making a tool, borrowing materials from a friend, etc.) may be used in another unique sequence at a later time. Therefore, effort, fitness token-payoffs (rewards), opportunity costs, risks, and many other components of evaluation need to be assigned continually to classes of acts. For this reason, there should be mechanisms that assign hedonic and other motivationally informative values to acts (e.g., "dangerous", "painful", "effort-consuming", "informative", "fun", "socially approved"), tallied as intermediate weights in decision processes. Our stream of actions and daily experiences will be affectively "colored" by the assignment of these hedonic values. If our psychological mechanisms were not using

present outcomes to assign a common internal currency of hedonic weights to classes of acts, there would be no function to suffering, joy, and so on. Emotion mode obviously impacts the assignment of hedonic values to acts.

Energy level, effort allocation, and mood: Overall metabolic budget will be regulated by emotion programs, as will specific allocations to various processes and facilitation or inhibition of specific activities. The effort that it takes to perform given tasks will shift accordingly, with things being easier or more effortful depending on how appropriate they are to the situation reflected by the emotion (Tooby & Cosmides, 1990a). Thus, fear will make it more difficult to attack an antagonist, whereas anger will make it easier. The confidence with which a situation has been identified (i.e., emotional clarity) should itself regulate the effortfulness of situation-appropriate activities. Confusion (itself an emotional state) should inhibit the expenditure of energy on costly behavioral responses and should motivate more information gathering and information analysis. Nesse (1990) has suggested that the function of mood is to reflect the propitiousness of the present environment for action, a hypothesis with many merits. We hypothesized (Tooby & Cosmides, 1990a) a similar function of mood, based on recognizing that the action-reward ratio of the environment is not a function of the environment alone, but an interaction between the structure the environment and the individual's present understanding of it. (By "understanding", we mean the correspondence between the structure of the environment, the structure of the algorithms, and the weightings and other information they use as parameters). The phenomenon that should regulate this aspect of mood is a perceived discrepancy between expected and actual payoff. The suspension of behavioral activity accompanied by very intense cognitive activity in depressed people looks like an effort to reconstruct models of the world so that future action can lead to payoffs, in part through stripping away previous valuations that led to unwelcome outcomes. Depression should be precipitated by (1) a heavy investment in a behavioral enterprise that was expected to lead to large payoffs that either failed to materialize or were not large enough to justify the investment; or (2) insufficient investment in maintaining a highly valued person or condition that was subsequently lost (possibly as a consequence); or (3) gradual recognition by situation-detectors that one's long-term pattern of effort and time expenditure has not led to a sufficient level of evolutionarily meaningful reward, when implicitly compared to alternative life paths (the condition of Dickens' Scrooge). Discrepancies between expected and actual payoff can occur in the other direction as well: Joy or precipitated surges of happiness are an emotion program that evolved to respond to the condition of an unexpectedly good outcome. They function to recalibrate previous value states which had led to underinvestment in or underexpectation for the successful activities or choices. Moreover, energy reserves that were being sequestered under one assumption about future prospects can be released given new, more accurate expectations about a more plentiful or advantageous future. Similarly, one can be informed of bad outcomes to choices not made: For example, one might find out that a company one almost invested in went bankrupt, or that the highway one almost took was snowed in. Information of this kind leads to an strengthening of the decision variables used (experienced as pleasure), which is sometimes mistaken for pleasure in the misfortune of others. Reciprocally, one can be informed of good outcomes to choices not made, which will be experienced as unpleasant.

Recalibrational emotions, evolved regulatory variables, and imagined experience: Information about outcomes is not equally spread throughout all points in time and all situations. Some situations are information-dense, full of ancestrally stable cues that reliably predicted the fitness consequences of certain decisions or revealed important variables (e.g., discovering who your father really is or how good a friend someone has been to you) and could therefore be used to alter weightings in decision rules.

Indeed, we expect that the architecture of the human mind is full of evolved variables whose function is to store summary magnitudes that are useful for regulating behavior and computation. These are not explicit concepts, representations, or goal states, but rather registers or indices that acquire their meaning by the evolved behavior-controlling and computation-controlling procedures that access them. Such regulatory variables may include measures of: how valuable to the individual a mate is, a child is, one's own life is, etc.; how stable or variable the food productivity of the habitat is; the distribution of condition-independent mortality in the habitat; one's expected future lifespan or period of efficacy; how good a friend someone has been to you; the extent of one's social support; one's aggressive formidability; one's sexual attractiveness; one's status or self-esteem; the status of the coalition one belongs to; present energy stores; present health; the degree to which subsistence requires collective action, and so on.

Most evolutionarily recurrent situations that select for emotion programs involve the discovery of information that allows the recomputation of one or more of these variables. Recalibration (which, when consciously accessible, appears to produce rich and distinct feeling states) is therefore a major functional component of most emotion programs. Jealousy, for example, involves several sets of recalibrations (e.g., diminution in estimate of own mate value, diminution of trust). Indeed, *recalibrational emotion programs* are emotion programs such as guilt, grief, depression, shame, and gratitude, whose primary function is to carry out such recomputations (Tooby & Cosmides, 1990a), rather than to orchestrate any short run behavioral response. These are emotion programs that have appeared puzzling from a functional perspective because the feelings they engender interfere with short-term utilitarian action that an active organism might be expected to engage in.

Consider guilt. Hamilton's (1964) rule defines the selection pressures that acted to build the circuits governing how organisms are motivated to allocate benefits between self and kin. This rule says nothing, however, about the procedures by which a mechanism could estimate the value of, say, a particular piece of food to oneself and one's kin. The fitness payoffs of such acts of assistance vary with circumstances. Consequently, each decision about where to allocate assistance depends on inferences about the relative weights of these variables. These nonconscious computations are subject to error. Imagine a mechanism that evolved to allocate food according to Hamilton's rule, situated (for example) in a hunter-gatherer woman. The mechanism in the woman has been using the best information available to her to weight the relative values of the meat to herself and her sister, perhaps reassuring her that it is safe to be away from her sister for awhile. The sudden discovery that her sister, since she was last contacted, has been starving and has become sick functions as an information-dense situation allowing the recalibration of the algorithms that weighted the relative values of the meat to self and sister. The sister's sickness functions as a cue that the previous allocation weighting was in error and that the variables need to be reweighted-including all of the weightings embedded in habitual action sequences. We believe that guilt functions as an emotion mode specialized for recalibration of regulatory variables that control trade-offs in welfare between self and other (Tooby & Cosmides, 1990a).

One significant subcomponent of these recomputational bouts is imagined experience, including both factual and counterfactual elements to potentiate branching decision points and the variables that govern them (Cosmides & Tooby, in press). Previous courses of action are brought to mind (I could have helped then, why didn't I think to?), with the effect of resetting choice points in decision rules. The negative valence of depression may be explained similarly: former actions that seemed pleasurable in the past, but which ultimately turned out to lead to bad outcomes, are re-experienced in imagination with a new affective coloration, so that in the future entirely different weightings are called up during choices.

Recalibrational releasing engines: The environment of evolutionary adaptedness was full of event relationships (e.g., mother is dead) and psychophysical regularities (e.g., blood indicates injury) that cued reliable information about the functional meanings and properties of things, events, persons, and regulatory variables to the psychological architecture. For example, certain body proportions and motions indicated immaturity and need, activating the emotion program of experiencing cuteness (see, Eibl-Eibesfeldt, 1970). Others indicated sexual attractiveness (Symons, 1979; Buss 1994). To be moved with gratitude, to be glad to be home, to see someone desperately pleading, to hold one's newborn baby in one's arms for the first time, to see a family member leave on a long trip, to encounter someone desperate with hunger, to hear one's baby cry with distress, to be warm while it is storming outside: these all *mean* something to us. How does this happen? In addition to the situation-detecting algorithms associated with major emotion programs such as fear, anger, or jealousy, we believe that humans have a far larger set of evolved specializations that we call *recalibrational releasing engines* that involve situation-detecting algorithms, and whose function is to trigger appropriate recalibrations, including affective recalibrations, when certain evolutionarily recognizable situations are encountered. By coordinating the mental contents of individuals in the same situation (since both intuitively know that, for example, the loss of one's mother is, as a default, experienced as a sad and painful event), these programs also facilitate communication and culture-learning, both of which depend on a shared frame of reference. Although these pervasive micro-programs construct a great deal of our world, investigations are only beginning into adaptations of this nature.

The role of imagery and emotion in planning: Imagery is the representation of perceptual information in a format that resembles actual perceptual input. In the evolution of animal nervous systems, simpler designs preceded more complex designs. The evolutionary designs of all modern species, including humans, use distinctive constellations of perceptual inputs as signals of states of affairs (for the rabbit, the outline of a hawk silhouette means a hawk is swooping in). Consequently, the key to unlocking and activating many complex evolved decision and evaluation programs was chained to the present: to being in an environment displaying specific perceptually detectable cues and cue constellations (sweetness, predators, running sores, emotion expressions).

There is a large inventory of wisdom stored in such programs, but this information, initially, could only be used by organisms in the environment displaying the activating cues - a profound limitation. An important design advance was achieved when psychological architectures evolved in which these programs could be accessed by feeding a decoupled fictional or counterfactual set of perceptual images, or event-relations, so that the response of these programs could be experienced and analyzed as part of planning, and other motivational and recalibrational functions (Tooby & Cosmides, 1990a; Cosmides & Tooby, in press). For example, the earlier design would go into a fear emotion mode, and flee the predator when encountered. The new design could imagine that a planned course of action would, as a side-effect, bring it into confrontation with a predator, experience (in appropriately attenuated and decoupled form) the fear program, and recognize that prospective, potential course of action as one to be avoided.

Recreating cues through imagery in a decoupled mode triggers the same emotion programs (minus their behavioral manifestations), and allows the planning function to evaluate imagined situations by using the same circuits that evaluate real situations. This would allow alternative courses of action to be evaluated in a way similar to the way in which experienced situations are evaluated. In other words, image-based representations may serve to unlock, for the purposes of planning, the same evolved mechanisms that are triggered by an actual encounter with a situation displaying the imagined perceptual and situational cues. For example, imagining the death of your child can call up the emotion state you would experience had this actually happened, activating previously dormant algorithms and making new information available to many different mechanisms. As many have recognized, this simulation process can help in making decisions about future plans: Even though you have never actually experienced the death of a child, for example, an imagined death may activate an image-based representation of extremely negative proprioceptive cues that "tell" the planning function that this is a situation to be avoided. Paradoxically, grief provoked by death may be a byproduct of mechanisms designed to take imagined situations as input: it may be intense so that, if triggered by imagination in advance, it is properly deterrent. Alternatively-or additionally-grief may be intense in order to recalibrate weightings in the decision rules that governed choices prior to the death. If your child died because you made an incorrect choice (and given the absence of a controlled study with alternative realities, a bad outcome always raises the probability that you made an incorrect choice), then experiencing grief will recalibrate you for subsequent choices. Death may involve guilt, grief, and depression because of the problem of recalibration of weights on courses of action. One may be haunted by guilt, meaning that courses of action retrospectively judged to be erroneous may be replayed in imagination over and over again, until the reweighting is accomplished. Similarly, joyful experiences may be savored, that is, replayed with attention to all of the details of the experience, so that every step of the course of action can be colored with positive weightings as it is rehearsed, again, until the simulated experience of these pseudo- "learning trials" has sufficiently reweighted the decision rules. The same principle might explain why rape victims often report experiencing horrifying unbidden images of the attack for 6-18 months after it has happened: the mind is replaying the trauma, running it through various decision rules and inference procedures, sifting it for clues of how to avoid such situations in the future, giving a different affective coloration to some of the locations, behaviors, and decisions that preceded the attack, and connecting them to a weighting of just how bad the consequent outcome was. After the 6-18 month period, the unbidden images suddenly stop, in a way that is sometimes described as "like a fever breaking": this would be the point at which the calibration is either done or there is no more to be learned from the experience (on unbidden images after trauma, see Horowitz, 1978, p. 142). One might expect the same phenomenon in combat veterans, with post-traumatic stress syndrome being an extreme version in which, for some reason, the shut-off mechanism malfunctions (Pitman & Orr, 1995).

Culture, ontogeny, and individual differences: How this theory of emotion can be integrated with models

of culture, models of human development, and models of individual differences must be treated elsewhere (see Tooby & Cosmides, 1990b, for an extended analysis of the relationship between emotions and individual differences; Tooby & Cosmides, 1992 for a discussion of culture). It is important to recognize, however, that the claim that evolved emotion programs are reliably developing aspects of a universal human nature does not necessarily imply fixed and uniform outcomes either for individuals or for cultures. Computational programs often have large numbers of open parameters, allowing their expression in adults to be highly variable, and until the mapping of the emotion programs is done, and tested cross-culturally (as Ekman and his associates did for facial expression), the range of variation will not be known.

Heuristic functions of the theory

The discussion so far should give some indications of how this theoretical approach allows the construction of testable, functional models for each emotion, and for the relations between emotion programs and other aspects of psychological functioning. The existence of such a theory also allows the discovery of previously unsuspected emotion states. Consideration of recurrent situations our ancestors would have had to be good at solving can prompt one to look for emotion modes even if one has never experienced them oneself (Cosmides & Tooby, 1994). A possible example is hunting. Humans are not just prey, equipped with fear emotions: they have also been predators for millions of years. A hunting emotion mode (predatoriness) may involve a special state of alert attention; suppression of any desire to talk (even before a particular animal is being stalked); heightened ability to read the minds of your companions; heightened sense of hearing; activation of abilities to make inferences about the presence, mental states, and activities of prey.

Moreover, the functional definition of emotion given here invites the possibility that many well-known mental states should be recognized as emotion states, such as the malaise engendered by infectious illness, coma, shock, the appreciation of beauty, homesickness, sexual arousal, confusion, nausea, and so on. For example, when you are sick, initiating actions and going about your daily activities is more effortful than usual: your impulse is to stay home and lie still. Although you feel as if your energy reserves are depleted, at a physical level the same fat reserves and digestively delivered glucose are available. Malaise is a computational state, not a physical one, and is designed to cope with the adaptive problem of illness, shunting energy from behavior to the immune system and possibly signaling the need for aid. Similarly, when situation-detecting algorithms detect the presence of a very grave internal injury, or the potential for one as indicated by a major blow, these may trigger a mode of operation of the cognitive system that is designed to prevent *any* discretionary movement. The function of this, in a world before hospitals, was to prevent further injury from being done, minimize blood loss and internal hemorrhaging, and allow the mobilization of the body's resources towards repair of immediate threats to the life: coma. Note that a coma is not a physically mandated state of paralysis: it is a computational state (technically, "a state of unconsciousness from which the patient cannot be roused" (Miller, 1976, p. 46), or "unarousable unresponsiveness" (Berkow, 1992, p. 1398): It occurs even when there has been no damage to the motor system.

Puzzles of consciousness and phenomenology

Emotions have a species-typical computational design, even if the quality of people's conscious experience in an emotion state varies. Phenomena such as hypnotic blindness and blindsight - where people lack the conscious experience of seeing, yet can be shown to be processing visual information - demonstrate that a computational state can exist without a person being aware of it. Moreover, there are many double dissociations between awareness and physiological states. That amputees experience phantom limbs shows that one can be aware of a non-existent physiological state (such as the presence of a non-existent leg!), whereas anosagnosics are unaware of having a paralyzed limb, and deny it is true, even in the face of evidence (Prigatano & Schacter 1991). Phenomena such as these show that whether a person becomes aware of an internal state is governed by machinery quite separate from that which creates the state itself. Hence awareness of a state such as an emotion cannot be what defines the presence of that state. The theoretical approach to the emotions herein provides criteria for assessing whether someone is in an emotion state (i.e., is running a particular emotion program), regardless of whether they admit it or are aware of it (or whether their culture has a word for it). The study of emotion can coexist with individual differences in the extent to which people metacognize about, or otherwise become aware of, their own emotion states (see, e.g.,

Weinberger, 1990 on repressors). At present, there is no validated, widely agreed on theory of the nature or function of consciousness. While an eventual scientific understanding of consciousness will be an important breakthrough, the study of the emotions can proceed without becoming entangled in the limitations of our present lack of understanding of consciousness.

[Comments about this article? Leave a note in our guest book!](#)

References

Barkow, J., Cosmides, L. & Tooby, J. (Eds.). *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.

Berkow, R. ed.) 1992. *The Merck Manual of Diagnosis and Therapy, Sixteenth edition*. Rahway, NJ: Merck.

Buss, D. M. 1989. 1994. *The evolution of desire*. New York: Basic Books.

Cannon, W. (1929). *Bodily changes in pain, hunger, fear and rage. Researches into the function of emotional excitement*. NY: Harper & Row.

Cosmides, L. 1989. The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition* 31: 187-276.

Cosmides, L. & Tooby, J. 1987. From evolution to behavior: Evolutionary psychology as the missing link. In *The latest on the best: Essays on evolution and optimality*, J. Dupre (Ed.). Cambridge, MA: MIT Press.

Cosmides, L. & Tooby, J. (1997). Dissecting the computational architecture of social inference mechanisms. In: *Characterizing human psychological adaptations* (Ciba Symposium Volume #208). Chichester: Wiley.

Cosmides, L. & Tooby, J. in press. Consider the source: the evolution of cognitive adaptations for metarepresentation. In: *Proceedings of the Xth Annual Cognitive Science Symposium*, D. Sperber, (Ed.)

Daly, M. & Wilson, M. 1988. *Homicide*. New York: Aldine.

Daly, M., Wilson, M., & Weghorst, S. J. 1982. Male sexual jealousy. *Ethology and Sociobiology* 3: 11-27.

Derryberry, D. & Tucker, D. 1994. Motivating the focus of attention. In P. M. Neidenthal & S. Kitayama (Eds.), *The heart's eye: Emotional influences in perception and attention* (pp. 167-196). Diego, CA: Academic Press.

Eibl-Eibesfeldt, I. 1970. *Ethology: The biology of behavior*. New York: Holt, Reinhart & Winston.

Ekman, P. (Ed.) 1982. *Emotion in the human face. Second Edition*. Cambridge: Cambridge University Press.

Ekman, P. (1985). *Telling lies*. NY: Norton.

Ekman, P. 1994. Strong evidence for universals in facial expressions. *Psychological Bulletin*. 115: 268-87.

Ekman, P., Levenson, R. & Friesen, W. 1983. Autonomic nervous system activities distinguishes between emotions. *Science* 221: 1208-10.

Frank, R. 1988. *Passions within reason: The strategic role of the emotions*. NY: Norton.

- Fredrickson, B. 1998. What good are positive emotions? *Review of General Psychology*, 2, 300-319.
- Fridlund, A. (1994). *Human facial expression: An evolutionary view*. San Diego: Academic Press.
- Gallistel, C.R. 1995. In *The cognitive neurosciences*, M. S. Gazzaniga (Ed.). Cambridge, MA: MIT Press.
- Garcia, J. (1990). Learning without memory. *J. Cognitive Neuroscience*, 2, 287-305.
- Gazzaniga, M.S. (Ed.). 1995 / 2000. *The cognitive neurosciences*, M. S. Gazzaniga (Ed.). Cambridge: MIT Press.
- Gould, S.J. 1997. Darwinian Fundamentalism. *New York Review of Books*, June 12, 1997.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology* 7, 1-52.
- Hirschfeld, L. & Gelman, S. (Eds.). (1994). *Mapping the mind: Domain specificity in cognition and culture*. NY: Cambridge University Press.
- Hirshleifer, J. (1987). On the emotions as guarantors of threats and promises. In *The latest on the best: Essays on evolution and optimality*, J. Dupre (Ed.). Cambridge, MA: MIT Press.
- Horowitz, M. (1978). *Image formation and cognition*. NY: Appleton-Century-Crofts.
- Isen, A. 1987. Positive affect, cognitive processes, and social behavior. *Advances in Experimental Social Psychology*, 20, 203-253.
- Jackendoff, R. 1987. *Consciousness and the computational mind*. Cambridge, MA: MIT Press.
- Lazarus, R. & Lazarus, B. 1994. *Passion and reason*. New York: Oxford University Press.
- LeDoux, J. 1995. In search of an emotional system in the brain: Leaping from fear to emotion to consciousness. In *The cognitive neurosciences*, M. S. Gazzaniga (Ed.). Cambridge, MA: MIT Press.
- Mackie, D. & Worth, L. (1991). Feeling good, but not thinking straight: The impact of positive mood on persuasion. In J. Forgas (Ed.), *Emotion and social judgments* (pp. 201-219). Oxford: Pergamon.
- Marks, I. (1987). *Fears, phobias, and rituals*. NY: Oxford.
- Marr, D. 1982. *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman.
- Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge University Press.
- Miller, S. (ed.) 1976. *Symptoms: The complete home medical encyclopedia*. New York: Crowell.
- Mineka, S. & Cook, M. (1993). Mechanisms involved in the observational conditioning of fear. *Journal of Experimental Psychology: General*, 122, 23-38.
- Nesse, R. (1990). Evolutionary explanations of emotions. *Human Nature*, 1, 261-289.
- Pinker, S. 1997. *How the mind works*. New York: Norton.
- Pitman, R. & Orr, S. (1995). Psychophysiology of emotional and memory networks in posttraumatic stress disorder. In J. McGaugh, N. Weinberger, & G. Lynch (Eds.), *Brain and memory: Modulation and mediation of neuroplasticity* (p. 75-83). NY: Oxford.

- Prigatano, G. & Schacter, D. (1991). *Awareness of deficit after brain injury*. NY: Oxford.
- Plutchik, R. 1980. *Emotion: A psychoevolutionary synthesis*. New York: Harper & Row.
- Raymond, (1991). *The New Hacker's Dictionary*. Cambridge, MA: MIT Press.
- Shepard, R. N. 1987. Evolution of a mesh between principles of the mind and regularities of the world. In *The latest on the best: Essays on evolution and Optimality*, J. Dupre (Ed.). Cambridge, MA: The MIT Press.
- Shepher, J. 1983. *Incest: a biosocial view*. New York: Academic Press.
- Swets, J. A., Tanner, W. D., & Birdsall, T. G. 1964. Decision processes in perception. In *Signal detection and recognition by human observers*, J. A. Swets (Ed.). Wiley.
- Symons, D. 1979. *The evolution of human sexuality*. Oxford: Oxford University Press.
- Tomaka, J., Blascovich, J., Kibler, J., Ernst, J. (1997). Cognitive and physiological antecedents of threat and challenge appraisal. *Journal of Personality & Social Psychology*, 73, 63-72.
- Tooby, J. 1985. The emergence of evolutionary psychology. In *Emerging Syntheses in Science*. Proceedings of the Founding Workshops of the Santa Fe Institute. Santa Fe, NM: The Santa Fe Institute.
- Tooby, J. & Cosmides, L. 1990a. The past explains the present: emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology*. 11: 375-424.
- Tooby, J. & Cosmides, L. 1990b. On the universality of human nature and the uniqueness of the individual: The role of genetics and adaptation. *Journal of Personality*. 58:
- Tooby, J. & Cosmides, L. 1992. The psychological foundations of culture. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind*. New York: Oxford University Press.
- Tooby, J. & Cosmides, L. 1996. The computationalist theory of communication. Paper given at the Human Behavior and Evolution Society Meetings, University of Arizona
- Tooby, J. & Cosmides, L. under review. On evolutionary psychology and modern adaptationism: A reply to Stephen Jay Gould.
- Williams, G. C. 1966. *Adaptation and natural selection: A critique of some current evolutionary thought*. Princeton: Princeton University Press.