

Baltic International Yearbook of Cognition, Logic and Communication

Volume 7 *MORALITY AND THE COGNITIVE SCIENCES*

Article 4

2012

Evolution and Moral Diversity

Tim Dean

University of New South Wales

Follow this and additional works at: <https://newprairiepress.org/byclc>



This work is licensed under a [Creative Commons Attribution-Noncommercial-No Derivative Works 4.0 License](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Recommended Citation

Dean, Tim (2012) "Evolution and Moral Diversity," *Baltic International Yearbook of Cognition, Logic and Communication*: Vol. 7. <https://doi.org/10.4148/byclc.v7i0.1775>

This Proceeding of the Symposium for Cognition, Logic and Communication is brought to you for free and open access by the Conferences at New Prairie Press. It has been accepted for inclusion in *Baltic International Yearbook of Cognition, Logic and Communication* by an authorized administrator of New Prairie Press. For more information, please contact cads@k-state.edu.

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2012
pages 1-16

Volume 7: *Morality and the Cognitive Sciences*
DOI: 10.4148/biyclc.v7i0.1775

TIM DEAN
University of New South Wales

EVOLUTION AND MORAL DIVERSITY

ABSTRACT: If humans have an evolved moral psychology, then we should not expect it to function in an identical way between individuals. Instead, we should expect a diversity in the function of our moral psychology between individuals that varies along genetic lines, and a corresponding diversity of moral attitudes and moral judgements that emerge from it. This is because there was no one psychological type that would reliably produce adaptive social behaviour in the highly heterogeneous environments in which our minds evolved. As such, there was no single psychological type towards which evolution could gravitate. Instead we evolved a stable polymorphism of psychological types, maintained by frequency-dependent selection, each predisposed towards particular social and moral attitudes and behaviours. This can help explain the existence of moral diversity, particularly intracultural moral diversity, and seemingly intractable moral disagreement.

One of the greatest adaptive challenges we faced as a species was how to live successfully in large social groups of unrelated individuals, thus reaping the benefits of cooperation and coordinated endeavour. However, despite its benefits, cooperation is a risky practice, particularly because it exposes individuals to the ill consequences of free-riders and defectors. The profound benefits of such coordinated group activity on individual fitness appears to have driven the remarkable

evolution of our species' cognitive capacities, even more so than the challenges of solving problems of engaging with our external environment (Sterelny 2012). To this end, there is mounting evidence that *homo sapiens* have evolved a slew of psychological faculties that are geared towards promoting prosocial behaviour in ourselves and others, including (but not limited to) the 'moral' emotions (Greene & Haidt 2002; Haidt 2001; Huebner et al. 2009), heuristics such as incest avoidance (Lieberman et al. 2003; Lieberman 2008), social imitation and in-group conformity (McElreath et al. 2003), and a tendency to create and adhere to behavioural norms (Sripada 2005). In concert, these mechanisms are capable of overcoming the substantial evolutionary hurdles to maintaining prosocial behaviour and cooperation in unrelated groups (Kitcher 2011).

However, while there has been tremendous progress over the past decade in revealing the psychological mechanisms that influence our moral attitudes, and their possible evolutionary influences, there has been relatively little attention given to the notion that individual differences contribute to variation in moral attitudes and behaviour. Most moral psychology studies tend to seek universal patterns that underlie moral judgements formed by normal functioning individuals within a particular environment. The research that does investigate the impact of individual difference on moral judgement had tended to focus on psychopathy or brain damage (Casebeer & Churchland 2003; Huebner et al. 2009). Much evolutionary psychology also concerns itself with uncovering the universals of human nature, the cognitive modules that have been shaped by natural selection over generations and which reside in us all (Cosmides & Tooby 2004). Relatively little evolutionary psychology literature addresses the influence of evolution on individual differences in psychology and cognitive function—although there are exceptions (Bateson 2004; Buss 2009)—and even less on the influence that individual differences have on moral diversity and moral disagreement.

While some have offered explanations of moral diversity between cultures within the existing framework of moral psychology research, such as in terms of the internalisation of different rules (Gill & Nichols 2008), which are established and fixed through punishment (Sripada 2005), along with the diversifying tendency to send clear signals of in-

group and out-group membership (Richerdson & Boyd 2001), explaining moral diversity *within* cultures is a significantly more difficult task. Explaining the source of these intra-cultural disagreements ought to be of great interest not only to moral psychologists seeking a comprehensive picture of how individuals form the moral attitudes and judgements they do, but also to moral philosophers and metaethicists. Even if it turns out that moral realists and moral universalists are correct that there is one true moral system, an understanding of what contributes to moral disagreement will still potentially be highly useful in working out how to dissolve those disagreements. Such realist and universalist philosophers might regard moral disagreement as an artefact of ignorance of the salient facts (moral or otherwise), cognitive impairment or bias, rather than being a genuine case of moral pluralism (Bloomfield 2001), but they still have to explain how moral disagreement emerges, and how to bring the various deluded interlocutors around to the ‘correct’ view. Explaining how moral diversity emerges is also of interest to moral anti-realists and subjectivists, who might see the diversity as telling us something useful about how humans construct morality and how these constructions disagree.

In this paper I argue for two main points: firstly, that a substantial amount of intra-cultural moral diversity can be attributed to individual differences in the operation of our moral psychology; secondly, that these individual differences exist because the very nature of the problems of social living meant that evolution was not able to settle upon a single psychological type that reliably produces adaptive behaviour in every social environment. Instead, a diversity—or ‘polymorphism’—of psychological types working together tended to be more evolutionarily stable, thus maintaining the polymorphism of psychological types in our species over time.

1. MORAL AND POLITICAL DIVERSITY

While there is a dearth of research in moral psychology directly tackling the question of the significance of individual difference in personality and cognitive function on moral proclivities, there is another source of literature that could act as a proxy and lend some insight into the impact of psychological variation on moral diversity: political psychology.

This field, which has undergone something of an explosion in growth over the last decade, is primarily concerned with understanding why individuals identify with certain political ideologies rather than others, and how they come to form the political attitudes they do. While traditional political science takes a top-down view of ideology, where ideologies are developed by political elites and disseminated to the public with varying degrees of accuracy (Zaller 1992), political psychology concerns itself with the bottom-up processes, often unconscious, by which people come to identify with an ideology, even if they lack a sophisticated appreciation of what the ideology entails according to the top-down account of political attitude formation (Thorisdottir et al. 2009). The analogy with moral psychology is clear, with both teasing out the psychological processes that lead to the formation of normative attitudes, although in different, but not dissimilar, domains. However, political psychology benefits from the assumption there is already diversity in attitudes, and seeks not to explain such diversity away, but rather explain why such diversity exists.

Another compelling link between political and moral psychology is that many of the ideological attitudes being assessed in political psychology have a moral dimension, such as attitudes towards equality, fairness, punishment, the treatment of outsiders or individuals of a different socioeconomic class, attitudes towards war, and even highly morally charged issues such as attitudes towards abortion or euthanasia. Some moral psychologists have already drawn a link between moral attitudes and political attitudes, with Jonathan Haidt and Jesse Graham finding that self-identifying liberals and conservatives rate certain moral issues—such as those concerning harm compared to those concerning authority—with a different level of importance, with liberals rating harm/care and fairness/reciprocity as more important than authority/respect, in-group/loyalty and purity/sanctity, and conservatives rating them all as similarly important (Haidt & Graham 2007).

However, many other political psychology studies have found variation in political attitudes to be significantly associated with other personality variables. One example is associating the psychological need to manage uncertainty and threat—real or imagined—with politically conservative attitudes (Jost et al. 2007). This study found that while everyone is motivated to minimise uncertainty in their lives to some ex-

tent, there is considerable variation in the extent to which individuals experience uncertainty as aversive and how they choose to resolve the uncertainty. There is similar variation in tolerance of ambiguity and its opposite, a tendency to stick to dichotomous conceptions and hold attitudes dogmatically (Frenkel-Brunswik 1948). Another variable is the propensity to perceive the world as being a dangerous place, and to have a somewhat pessimistic assumption about human nature at large (Altemeyer 1998; Duckitt 2001). All¹ the above are strongly correlated with holding politically conservative attitudes.

Another aspect of psychology that has been studied is integrative complexity, which is termed a ‘cognitive style,’ and refers to how people tend to integrate and process information. An individual with low integrative complexity will tend to take a black and white view of issues and will employ simple evaluative categories to attitudes, such as ‘good’ or ‘bad’ rather than taking a more complex view considering strengths and weaknesses of one particular notion (Tetlock 1983). High integrative complexity is correlated with liberal views, while low integrative complexity is associated with conservative views.²

One source of variation that would, on the surface, appear to have little to do with politics is personality as measured by the Five Factor Model (the ‘Big Five’): openness, conscientiousness, extraversion, agreeableness and neuroticism. However, two of these variables have considerable predictive power when it comes to political attitudes. High levels of openness—which is often defined in terms of a propensity to seek out novel stimuli, to engage in intellectual pursuits, to be creative and express a general intellectual curiosity—are positively correlated with liberal self-identification and liberal views. High levels of conscientiousness—which is often associated with being organised, dependable, punctual and self-controlled—are positively correlated with conservative self-identification and attitudes. The other three factors have either conflicting or inconclusive evidence suggesting correlation with political attitudes (Mondak 2010). The finding that personality is somehow predictive of political attitudes is particularly illuminating, because personality is something that appears to be relatively stable over one’s life, and is robustly heritable (McCrae & Costa 2003), suggesting a strong genetic component. This is also borne out by studies that find political attitudes are also strongly heritable (Alford et al.

2005), a notion to which I will return below.

As for the mechanism that links psychology and politics, it appears as though one’s ‘worldview’—their broad framework for perceiving and understanding the world, and investing it with meaning and value—has a heavy influence on political attitudes (Lakoff 1996). So if someone perceives the world to be a dangerous place, they are more likely to hold conservative views than someone who perceives the world to be a relatively safe place (Jost et al. 2003). Likewise, if they perceive the world to be a meritocracy, where people reliably get what they deserve—both rewards and punishments—then they are more likely to tilt towards conservative attitudes (McCoy & Major 2007). This phenomenon also carries over to worlds constructed in the imagination via a thought experiment, with individuals self-identifying as either liberals or conservatives exhibiting a noticeable tilt towards conservative attitudes when the world is presented as being either dangerous or a meritocracy (Mitchell & Tetlock 2009). Yet, when the evidence as to whether the world is dangerous or a meritocracy is ambiguous, variation in attitudes re-emerges, and does so along the lines of the psychological variations mentioned above.

This suggests that the link between psychology and politics is mediated by a middle step—one’s ‘worldview’—such that an individual’s worldview has a strong influence on their political attitudes. Yet, their personality and other cognitive proclivities have a strong influence on their worldview. For example, someone who has a naturally strong fear response, or a naturally low tolerance for ambiguity, is likely to undergo a very different experience when presented with a certain situation than an individual who has a mild fear response and high tolerance for ambiguity. As such, the very same environment can alter one’s worldview dependent on one’s experience of it. This, in turn, makes an individual more likely to find particular political ideologies and attitudes more appealing, such as someone with a strong fear response being more attracted to the more security-inclined approach offered by conservatism. However, ideology also works to influence worldview, leading to a feedback that serves to reinforce the worldview associated with a particular ideology. The strength of this effect, if it exists, is unknown, and will likely remain so until specific empirical studies are conducted to put it to the test.

There is not only evidence for a link between individual differences in psychology and political attitudes, but also for a link between individual differences in psychology and genetics (McCrae & Costa 2003). Indeed, there is evidence that genes extend their influence all the way to political attitudes (Alford et al. 2005). While there does not appear to be any single gene or cluster of genes for conservatism or liberalism—and it is unlikely that one will ever be found—it does appear as though some people are genetically predisposed towards either liberal or conservative attitudes. The question that stems from this curious finding is why might biology predispose people towards one set of political attitudes or another? The answer may come from that great synthesiser: evolution.

2. EVOLUTION OF SOCIAL COGNITION

It appears the main driver of human cognitive evolution over the past few million years was the tremendous complexity of social dynamics, which in turn was due to the huge adaptive benefits of social and cooperative behaviour (Sterelny 2007). The greater social complexity placed elevating demands on our cognitive faculties, even more so than the demands of adapting to the physical environment. After all, directing behaviour in a highly heterogeneous environment is far more complex a task than steering behaviour in a relatively homogeneous environment, and the social landscape is orders of magnitude more heterogeneous than the physical landscape. It is one thing to track and respond to the relatively static physical world around you, another entirely to track a multitude of relationships within one's social cluster and estimate the relative status of those involved, track their intentions, predict their behaviour and their reactions to one's own, and absorb and adapt to the culturally transmitted information passed between individuals. Greater social complexity also enabled more ecological innovation, thus raising the heterogeneity of the ecological landscape to be navigated by early humans, triggering a feedback loop that further increased cognitive demands (Sterelny 2007). As a result, humans evolved a slew of highly sophisticated cognitive mechanisms to enable us to successfully steer behaviour, particularly in the highly heterogeneous social landscape, including many mecha-

nisms that are presently the focus of moral psychology studies, such as moral emotions like empathy and guilt (Haidt 2003), heuristics such as social imitation and in-group conformity (Richerdson et al. 2002), and a tendency to create and adhere to behavioural norms (Sripada 2005).

One cunning way of dealing with a heterogeneous environment is phenotypic plasticity—i.e. a single genotype can result in multiple different phenotypes, each adaptive in a particular environment—and there is ample evidence that the human mind is tremendously plastic for this very reason (Godfrey-Smith 1998). A plastic mind can adapt to its environment by innovating new behavioural strategies ‘on the fly’ rather than waiting for the cumbersome process of innovation through genetic mutation filtered by natural selection to fix new behavioural tendencies over the course of generations. However, plasticity has its downsides. There is a cost to plasticity, particularly when it comes to the concomitant faculties required to accurately track the state of the environment and steer the phenotype to be adaptive in that environment. Adaptive environments are rarely fully transparent, and this is particularly the case with social environments, where some individuals will be inclined to send false signals in order to manipulate the behaviour of others to their advantage. Another cost of plasticity is that it takes time and energy to track the environment, innovate a new behavioural strategy to deal with it, and then employ it. An organism with a relatively low cost fixed phenotype will likely prove more adaptive than an organism with more costly plasticity in environments where the fixed phenotype is consistently adaptive. If the environment turns out to be relatively homogeneous, then the plastic organism has a relative disadvantage.

Another limit to plasticity is that plasticity is theoretically limitless; no organism can be infinitely plastic. The question is: how plastic is optimal? Given the above trade-off between fixed and plastic phenotypes, the optimal amount will likely depend on how heterogeneous the environment really is. However, many environments won't be evenly heterogeneous. This might be particularly true of social environments. Some aspects might remain fairly static, while others might change dramatically over a short time. Some features of the social environment will vary depending on where the individual is born, but once detected, will remain fixed for life. Others will vary throughout the

lifetime of the organism. The physical environment might also change over time, albeit likely at a slower rate. The upshot is that there is not only heterogeneity to deal with, but a *heterogeneity of heterogeneity*, which I call H^2 .

One way to deal with H^2 is to constrain plasticity in some aspects while employing it in others. One way would be to limit plasticity to very specific domains, and there is evidence that humans do just this, such as in language or the emotion of disgust being ubiquitous faculties but the specific language spoken or triggers for disgust being environmentally contingent. Another way would be to vary the amount of plasticity in the population. As such, some genotypes might be predisposed towards favouring particular phenotypes, while other genotypes will be predisposed towards favouring others, with considerable overlap and plasticity in between. Or there could be variation in how phenotypically or behaviourally innovative they are. Or variation in how likely they are to mimic the behaviours of others. Or variation in their predisposition towards certain behaviours over others.

This constrained plasticity appears to be precisely the solution that evolution stumbled upon when it comes to many of our psychological proclivities, such as personality. Personality represents the broad behavioural tendencies that individuals exhibit, while leaving considerable room for behavioural plasticity. And personality, while highly polymorphic, appears to be highly heritable (Bouchard 1994; Plomin et al. 1994; Turkheimer 2000; Bouchard & McGue 2003). The reason for this appears to be that the polymorphism of personality types is maintained by balancing selection in response to high environmental heterogeneity. This form of selection occurs when there is no single allele that proves sufficiently adaptive to achieve fixation within the population, and a polymorphism of alleles is maintained instead (Penke et al. 2007).

3. STRATEGIC DIVERSITY

The key to maintaining this polymorphism is the fact that different behavioural strategies yield different results³ depending on the environment in which they are employed. For example, the behavioural strategies that yield the best results in a resource-poor environment

may well be quite different from the behavioural strategies that succeed in a resource-rich environment. In the former, the costs of failure may be such that there is little scope for engaging in risky endeavours, such as a cooperative venture that leaves an individual exposed to mortal threat should their companion choose to free-ride on their efforts. However, if the costs of failure are lower in the resource-rich environment, then such cooperative endeavours might yield a greater outcome, even if they occasionally suffer from free-riding.

Unlike the physical environment, which remains fairly stable over time, the social environment is constantly shifting as individuals modify their behaviour to improve their payoffs in response to the behaviours of others. Adapting behavioural strategies to these two environments presents radically different problems. Adapting to the physical environment is relatively easy, at least compared to the social environment. A successful behavioural strategy such as ‘avoid drinking water with a foul odour’ or ‘be wary of long slithering creatures’ will likely remain successful in a particular physical environment for many generations, thus enabling the trait that predisposes this strategy to become fixed in the population. The physical environment also typically won’t react as quickly in response to the presence of such behavioural traits.

Adapting to the social environment is an entirely different matter. As new behavioural strategies are innovated, the population will adapt to them and develop new strategies in response, often quite rapidly. In this environment, the optimal strategy is often a moving target. One example from the Big Five personality traits (Openness, Conscientiousness, Extraversion, Agreeableness, Neuroticism) is the apparent trade-off involved in the introversion-extraversion spectrum. While extraverts might enjoy greater mating success than introverts, this benefit comes with a cost of increased physical risk and/or decreased parenting effort (Nettle 2005). As a result, the polymorphism in extraversion-introversion in a population is maintained by negative frequency-dependent selection, which favours rare alleles: too many extraverts in a population, and introversion becomes more adaptive, and vice versa.

Such frequency-dependent selection also appears to be at work when it comes to behavioural dispositions that affect cooperative be-

haviour. The dynamics of cooperative interaction have been extensively studied using the tools of game theory, and the dynamics of repeated cooperative interactions between multiple individuals have been very effectively modelled particularly by the Iterated Prisoner's Dilemma (Trivers 1971; Axelrod & Hamilton 1981; Binmore 1998). This game involves a population of individuals, or 'agents,' who repeatedly interact with other members of the population in one-on-one Prisoner's Dilemmas. The dynamics of the Prisoner's Dilemma are particularly apt at representing the problems of cooperative interaction, with the payoff for each individual dependent on the actions of the other. If both agents choose to cooperate, they both get a high payoff, called R for 'reward.' If one chooses to cooperate but the other chooses to defect by exploiting that cooperation without reciprocating, the defector receives an even higher payoff, called T for 'temptation,' while the hapless cooperator receives little or nothing in return, called 'S' for 'sucker's payoff.' Should they both defect—either in an attempt to exploit the other's cooperation or to avoid the sucker's payoff—then both get a low reward, called 'P' for punishment. The rewards are always constructed such that $T > R > P > S$, and $2R > T+S$.

One of the more notable findings from studies of the IPD is that it is almost impossible to sustain cooperation at maximal—or Pareto optimal (PO)—levels. The PO outcome would be for perpetual mutual cooperation, whereby each individual receives R, which is greater than temptation plus the sucker's payoff. However, an environment where the entire population is certain to cooperate is the very paragon of a defector's paradise. Assuming even a small chance of innovating new strategies, either randomly or intentionally, then it is virtually inevitable that a population will be 'invaded' by at least one defector. Should the payoffs represent reproductive fitness, then the hapless cooperators will soon be outcompeted by the defectors.

Even so, when the IPD simulations are constructed to approximate real-world conditions, such as assuming agents are not perfectly rational but are instead boundedly rational,⁴ they show that populations often do sustain levels of cooperation above that of mutual defection. Instead, they often evolve a *mix* of strategies after many iterations (Axelrod 1987). For example, some strategies are highly cooperative, while others are more cautious and are more prone to defect. This mix

of strategies interacts in such a way as to promote aggregate levels of cooperation within the group without being crippled by the deleterious effects of defection or free-riding. While the level of cooperation may not reach PO levels, it is certainly higher than the rock-bottom level of mutual defection. This dynamic emerges because there is no one Nash equilibrium in the IPD, meaning there is no one stable state consisting of any one pure⁵ strategy that cannot be successfully invaded by other strategies (Boyd & Lorberbaum 1987). The end result is a mix of strategies—or a polymorphism of types that produce certain strategies—each with its own merits given the social 'environment.' This results in a kind of dynamic equilibrium whereby the system refuses to settle into a single stable state. Instead the strategies in play are constantly changing in frequency, adapting to the current environment, but unable to find a state where any single strategy dominates the others. One particularly interesting such state consists of a pluralism of 'robust' strategies—strategies that perform well against many other strategies—that work synergistically to produce a 'core' of highly cooperative strategies surrounded by a protective perimeter 'shield' of suspicious strategies, which protect the core from invasion by 'nasty' strategies on the periphery (Lomborg, 1996).

Thus frequency-dependent selection could select for a polymorphism of psychological types that are disposed to produce a mix of behavioural strategies which, along with high plasticity, allow individuals within a population to produce adaptive behaviour in the H^2 social environment. As infinite plasticity is unfeasible, and there is no one psychological type that will reliably produce adaptive behaviour in every social environment, a polymorphism of psychological types is maintained, along with considerable plasticity. I suggest that it is this diversity of psychological types that contributes, in part, to the diversity in worldviews that individuals adopt, and this diversity of worldviews in turn contributes to diversity in their moral and political attitudes. Thus, at least some moral diversity and moral disagreement is caused by variation in psychological types, and produced by differences in the way individuals experience the world, which contributes to their differing attitudes and beliefs.

4. CONCLUSION

If it is the case that genes influence moral attitudes via personality and psychology, what does this imply for moral psychology and moral philosophy? While moral issues and political issues are often conceived of as belonging to different realms, many of the problems that political attitudes and moral attitudes are directed towards solving are very similar, whether it be problems of dealing with those who break the rules of social behaviour, of the permissibility or impermissibility of certain practices, such as abortion or euthanasia, or the fair distribution of wealth and opportunity within a society. The overlap is substantial enough that the results in political psychology warrant attention from moral psychologists. And the findings mentioned above suggest, at least on a descriptive level, that morality could be even more complex than we might have thought (or hoped). Moral attitudes might be produced by variations in psychological predispositions, driven by an evolved need to develop behavioural strategies that are successful in highly heterogeneous, and H^2 , social environments. This suggests that moral attitudes are best seen as constituting solutions to problems of social living. And to the extent that these problems have no single solution, or no single set of strategies that will prove successful in every environment, then perhaps such diversity in attitudes might prove expeditious at promoting social stability in a wide range of environmental conditions.

The next step will be to put this theorised association between psychological variation and moral diversity to the test. The research in political psychology is intriguing, but it is not asking the same questions as moral psychology. However, I strongly suspect that should moral psychologists develop the appropriate tests, they will find that variation in psychology and personality will contribute to variation in moral judgement and behaviour. We just need to go out and see.

Notes

¹The terms 'liberal' and 'conservative' here refer primarily to the United States political spectrum, although elements can be abstracted to other political cultures.

²An anecdotal example might be the difference between former United States President, George W. Bush, and the Democratic presidential candidate in 2004, Senator John

Kerry. The latter was criticised by conservatives and conservative elements of the media for being a "flip-flopper" because of his well known propensity to change his mind on particular policies. However, liberals praised this tendency, noting that he changed his mind on the strength of new evidence or argument. They, in turn, criticised George W. Bush for being overly simplistic and black and white in his thinking, a trait praised for its clarity and decisiveness by conservatives.

³The term "results" can be cashed out in a number of ways, such as by referring to payoffs for individuals in cooperative endeavours, such as the amount of food garnered for each individual in a cooperative hunt. However, the most important sense is that of evolutionary fitness for the individuals involved in the cooperative endeavour.

⁴By 'boundedly rational' I mean they are not presumed to be perfectly error-free rational agents that have access to all the relevant information and can execute strategies of arbitrary complexity. Boundedly rational agents can be modelled in many ways. For example, (Lomborg 1996) models his agents as using "simple rule-of-thumb strategies that work well in their current surroundings." He gives them limited memory, allows them to imitate successful strategies with which they interact, allows them to innovate new strategies and introduces noise into the simulation to represent errors in information exchange, judgement or action.

⁵A 'pure' strategy is one that exhaustively describes how an agent will act in any situation with fixed probabilities. Strategies such as 'always cooperate', 'always defect' and 'tit-for-tat' are pure strategies. A strategy that cooperates $x\%$ of the time is considered a mixed rather than pure strategy.

References

- Alford, J. R., Funk, C. L. & Hibbing, J. R. 2005. 'Are Political Orientations Genetically Transmitted?' *American Political Science Review* 99, no. 2: 153–167.
- Altemeyer, B. 1998. 'The Other Authoritarian Personality'. *Advances in experimental social psychology* 30: 47–92.
- Axelrod, R. 1987. 'The evolution of strategies in the iterated prisoner's dilemma'. *Genetic algorithms and simulated annealing* 3: 32–41.
- Axelrod, R. & Hamilton, W. D. 1981. 'The evolution of cooperation'. *Science* 211, no. 4489: 1390–1396.
- Bateson, P. 2004. 'The origins of human differences'. *Daedalus* 133, no. 4: 36–46.
- Binmore, K. G. 1998. *Game Theory and the Social Contract: Just Playing*. MIT Press.
- Bloomfield, P. 2001. *Moral reality*. Oxford University Press, USA.
- Bouchard, T. 1994. 'Genes, environment, and personality'. *Science* 264: 1700–1701.
- Bouchard, T. & McGue, M. 2003. 'Genetic and environmental influences on human psychological differences'. *Journal of neurobiology* 54, no. 1: 4–45.
- Boyd, R. & Lorberbaum, J. P. 1987. 'No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game'. Nature Publishing Group.
- Buss, D. M. 2009. 'How can evolutionary psychology successfully explain personality and individual differences?' *Perspectives on Psychological Science* 4, no. 4: 359–366.
- Casebeer, W. D. & Churchland, P. S. 2003. 'The Neural Mechanisms of Moral Cognition: A Multiple-Aspect Approach to Moral Judgment and Decision-Making'. *Biology and Philosophy* 18, no. 1: 169–194.

- Cosmides, L. & Tooby, J. 2004. 'Knowing thyself: The evolutionary psychology of moral reasoning and moral sentiments'. *Business, science, and ethics* 4: 91–127.
- Duckitt, J. 2001. 'A dual-process cognitive-motivational theory of ideology and prejudice'. *Advances in experimental social psychology* 33: 41–113.
- Frenkel-Brunswik, E. 1948. 'Tolerance toward ambiguity as a personality variable'. *American Psychologist* 3: 250.
- Gill, M. B. & Nichols, S. 2008. 'Sentimentalist Pluralism: Moral Psychology and Philosophical Ethics'. *Philosophical Issues* 18, no. 1: 143–163.
- Godfrey-Smith, P. 1998. *Complexity and the Function of Mind in Nature*. Cambridge University Press.
- Greene, J. & Haidt, J. 2002. 'How (and where) does moral judgment work?' *Trends in Cognitive Sciences* 6, no. 12: 517–523.
- Haidt, J. 2001. 'The emotional dog and its rational tail: A social intuitionist approach to moral judgment'. *Psychological Review* 108, no. 4: 814–834.
- . 2003. 'The moral emotions'. In R. J. Davidson, K. R. Scherer & H. H. Goldsmith (eds.) 'Handbook of affective sciences', 852–870. Oxford University Press, USA.
- Haidt, J. & Graham, J. 2007. 'Planet of the Durkheimians, where community, authority, and sacredness are foundations of morality'. *Social and psychological bases of ideology and system justification (December 11, 2006)* 371–401.
- Huebner, B., Dwyer, S. & Hauser, M. 2009. 'The role of emotion in moral psychology'. *Trends in cognitive sciences* 13, no. 1: 1–6.
- Jost, J. T., Glaser, J., Kruglanski, A. W. & Sulloway, F. J. 2003. 'Political conservatism as motivated social cognition'. *Psychological Bulletin* 129, no. 3: 339–375.
- Jost, J. T., Napier, J. L., Thorisdottir, H., Gosling, S. D., Palfai, T. P. & Ostafin, B. 2007. 'Are needs to manage uncertainty and threat associated with political conservatism or ideological extremity?' *Personality and social psychology bulletin* 33, no. 7: 989–1007.
- Kitcher, P. 2011. *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Lakoff, G. 1996. *Moral politics*. University of Chicago Press.
- Lieberman, D. 2008. 'Moral Sentiments Relating to Incest: Discerning Adaptations from By-products'. In W. Sinnott-Armstrong (ed.) 'Moral Psychology vol. 1.', 165–190. Cambridge, MA: The MIT Press.
- Lieberman, D., Tooby, J. & Cosmides, L. 2003. 'Does morality have a biological basis? An empirical test of the factors governing moral sentiments relating to incest.' *Proceedings. Biological sciences / The Royal Society* 270, no. 1517: 819–26.
- Lomborg, B. 1996. 'Nucleus and Shield: The Evolution of Social Structure in the Iterated Prisoner's Dilemma'. *American Sociological Review* 61, no. 2: 278–307.
- McCoy, S. & Major, B. 2007. 'Priming meritocracy and the psychological justification of inequality'. *Journal of Experimental Social Psychology* 43, no. 3: 341–351.
- McCrae, R. R. & Costa, P. T. 2003. *Personality in adulthood: A five-factor theory perspective*. The Guilford Press.
- McElreath, R., Boyd, R. & Richerds, P. J. 2003. 'Shared Norms and the Evolution of Ethnic Markers'. *Current Anthropology* 44, no. 1: 341–351.
- Mitchell, G. & Tetlock, P. E. 2009. 'Disentangling reasons and rationalizations: Exploring perceived fairness in hypothetical societies'. *Social and psychological bases of ideology and system justification* 1, no. 9: 126–158.
- Mondak, J. J. 2010. *Personality and the foundations of political behavior*. Cambridge Univ Pr.

- Nettle, D. 2005. 'An evolutionary approach to the extraversion continuum'. *Evolution and Human Behavior* 26, no. 4: 363–373.
- Penke, L., Denissen, J. J. A. & Miller, G. F. 2007. 'The evolutionary genetics of personality'. *European Journal of Personality* 21, no. 5: 549–587.
- Plomin, R., Owen, M. J. & McGuffin, P. 1994. 'The genetic basis of complex human behaviors'. *Science* 264, no. 5166: 1733–1739.
- Richerds, P. J. & Boyd, R. 2001. 'The evolution of subjective commitment to groups: A tribal instincts hypothesis'. In R. M. Nesse (ed.) 'The Evolution of Subjective Commitment', 186–220. Russell Sage Foundation Publications.
- Richerds, P. J., Boyd, R. & Henrich, J. 2002. 'Cultural Evolution of Human Cooperation'. In P. Hammerstein (ed.) 'Genetic and Cultural Evolution of Cooperation', 1–33. MIT Press.
- Sripada, C. S. 2005. 'Punishment and the strategic structure of moral systems'. *Biology and Philosophy* 20, no. 4: 767–789.
- Sripada, C. S. & Stich, S. 2005. 'A Framework for the Psychology of Norms'. In P. Carruthers, S. Laurence & S. Stich (eds.) 'The Innate Mind: Structure and Contents', 1–40. Oxford University Press, UK.
- Sterelny, K. 2007. 'Social intelligence, human intelligence and niche construction'. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences* 362, no. 1480: 719–30.
- . 2012. *The Evolved Apprentice*. Cambridge, MA: The MIT Press.
- Tetlock, P. E. 1983. 'Cognitive style and political ideology'. *Journal of Personality and Social Psychology* 45, no. 1: 118–126.
- Thorisdottir, H., Jost, J. T. & Kay, A. C. 2009. 'On the Social and Psychological Bases of Ideology and System Justification'. In J. T. Jost, A. C. Kay & H. Thorisdottir (eds.) 'Social and Psychological Bases of Ideology and System Justification', 1. Oxford University Press, USA.
- Trivers, R. L. 1971. 'The evolution of reciprocal altruism'. *The Quarterly review of biology* 46, no. 1: 35–57.
- Turkheimer, E. 2000. 'Three laws of behavior genetics and what they mean'. *Current Directions in Psychological Science* 9, no. 5: 160–164.
- Zaller, J. 1992. *The nature and origins of mass opinion*. Cambridge University Press.