Available online at www.sciencedirect.com
ScienceDirect
Cognition 107 (2008) 817–828
COGNITION
www.elsevier.com/locate/COGNIT

Enhanced recognition of defectors depends on their rarity

Pat Barclay *

Department of Neurobiology and Behavior, Cornell University, Ithaca, NY 14853, USA

Received 13 October 2006; revised 13 June 2007; accepted 30 November 2007

Abstract

Evolutionary psychologists have proposed that humans possess cognitive mechanisms for social exchange, but have perhaps focused overmuch on "cheating", because avoiding exploitation in reciprocal exchange could be accomplished either by avoidance of defectors or by attraction to cooperators. Past studies that have claimed to support the existence of a "cheater-detection module" by finding enhanced memory for the faces of "cheaters" have mostly relied on verbal descriptions, and these are prone to bias if the degree of cheating is unintentionally more severe than the degree of cooperation. Given that populations differ in the prevalence of defectors, it is most effective to remember whatever type is rare rather than always focus on cheaters. In the present experiment, participants played a computerized trust game and saw faces of cooperators and defectors in 20%/80%, 50%/50%, or 80%/20% ratios. Consistent with predictions, defectors were remembered best when rare but worst when common, supporting the existence of slightly more general cognitive mechanisms rather than specific cheater-recognition mechanisms.

Keywords: Cheater recognition; Domain specificity; Reciprocal altruism; Cooperation; Trust; Evolutionary psychology

* Tel.: +1 607 254 4377; fax: +1 607 254 1303.
E-mail address: pjb46@cornell.edu

0010-0277/$ - see front matter
doi:10.1016/j.cognition.2007.11.013

1. Introduction

When two individuals reciprocate generous acts towards each other, they can be better off than they would be without this cooperation (Axelrod & Hamilton, 1981; Trivers, 1971). For such reciprocal altruism to be stable, cooperators must avoid exploitation from defectors (those who do not cooperate/reciprocate). This requires cognitive abilities that can solve two tasks: (1) Detecting instances of non-cooperation (cheating detection), and (2) remembering who has been cooperative and who has not (cheater recognition) and interacting preferentially with other cooperators (Trivers, 1971). Cosmides and Tooby (e.g. 1992, 2000) have demonstrated that people are good at recognizing instances of cooperation and defection, while other researchers have used experimental games to show that people interact preferentially with cooperators and/or are more generous towards them (Barclay, 2004; Barclay & Willer, 2007; Fehr & Fischbacher, 2004; Komorita & Parks, 1995; Roberts & Renwick, 2003).

Many evolutionary psychologists have assumed that people accomplish these two tasks by focusing on cheaters and instances of non-reciprocation. For example, Cosmides and Tooby argue that humans possess a cognitive adaptation specifically designed for "cheater detection" (e.g. 1992, 2000), and have explicitly claimed that "detecting cheaters is necessary for contingent cooperation to evolve" (2000, p. 592, emphasis added). Others have proposed that people have a special memory for cheaters (e.g. Mealey, Daood, & Krage, 1996; Oda, 1997). Although it is certainly necessary for people to distinguish between instances of cooperation and defection (as argued by Cosmides, Tooby, and others), people could accomplish the task of remembering whom to cooperate with if they possessed cognitive mechanisms that caused them to focus on reciprocated altruism and to selectively remember and interact with cooperators rather than cheaters (Brown & Moore, 2000). From a selectionist point of view, it should not matter how people accomplish these tasks, so long as they are successfully accomplished with as low a cost and error rate as possible. In fact, the most successful psychology could focus attention in part on whatever action, cooperation or defection, is less common in a population and invest more in remembering instances of the rare action (and/or people who tend to do the rare action), and such a psychology could build on pre-existing tendencies to remember rarity (e.g. Hunt, 1995; Hunt, 2006; McDaniel & Geraci, 2006). There are many cooperative norms in our society (Fehr & Fischbacher, 2004), so cheating may seem more salient to us because of the relative prevalence of cooperation and rarity of deal-breaking. Past researchers may have focused on cheating because of this apparent salience. However, there are a few reasons to expect a more general cognitive system, perhaps one that focuses on reputation in general, to underlie social exchange rather than systems that focus specifically on cheaters.

First, such a flexible system could demand fewer cognitive resources than a psychology that always focuses on cheaters regardless of their frequency. By focusing on the rare type, one need not hold every single defector (or act of defection) in memory if it would be easier to remember the few rare cooperators in an uncooperative population, and this saves valuable and limited cognitive resources. When defectors are common, the cost of remembering numerous defectors can outweigh the costs of investing cognitive resources in remembering cooperators instead. This is especially crucial when reciprocal altruism first invades a population because there would be many more defectors than cooperators, and a cognitive mechanism designed to remember defectors would be at a disadvantage compared to one designed to remember cooperators or whichever type is rarer.

Secondly, a more general "reputation-tracking" system would likely commit fewer errors than a system that always focuses on cheaters. A "cheater-recognition module" would presumably use a decision rule like "remember the defectors and avoid them, assume everyone else you already know is a cooperator unless you remember them being a defector." When attempting to classify people as cooperators or defectors, one will make more misclassifications of whichever type is more common simply because there would be more of them to misclassify. Thus, a reciprocal altruist in a population of mostly defectors would make more errors if he/she tried to remember all of the defectors than if he/she tried to remember the few rare cooperators, whereas the reverse would be true in a population of mostly cooperators. One could have a decision rule along the lines of "be biased towards doing what most people do, unless your current partner stands out as one who will behave differently." Groups may differ in the amount of cooperation displayed (e.g. Henrich et al., 2005; Richerson & Boyd, 2005) and/or the number of defectors, so it could be more adaptive to have such a facultative mental module causing one to remember whichever type (cooperators/defectors) is rare in one's group, rather than a fixed mental module causing one to always focus on either cooperators or on defectors. If one is a member of multiple social groups, one might even track the proportion of defectors in each group and focus on whatever is rare in each group even though this may differ between groups.

Thirdly, such facultative memory could simply be a byproduct of a general tendency to focus on rarity and remember exceptions and atypical stimuli (e.g. Graesser, Woll, Kowalski, & Smith, 1980; Hunt, 2006). This principle of minimizing cognitive effort by investing effort in remembering rare types likely applies to the evolution of other cognitive mechanisms. Given that humans are naturally interested in others' actions (Barrett, Dunbar, & Lycett, 2002) and care about others' intentions (e.g. Nelson, 2002) and relative payoffs (e.g. Bolton & Ockenfels, 2000), this tendency to notice and remember rare types would naturally apply to social exchange, resulting in differential memory for cooperators or defectors. Alternately, this general tendency to notice rarity could have been a preadaptation that was later co-opted into a special mental design for social exchange, possibly by directing attentional resources towards others' reputation in general (which is then remembered by those more domain-general memory processes).

1.1. Past studies on cheater recognition

Several studies have examined whether people have better memory for "cheaters" than cooperators. In the best known of these, Mealey and colleagues (Mealey et al., 1996) presented faces alongside written descriptions of cheating, trustworthy, or neutral acts, and tested participants' memory for faces one week later. The resulting enhanced memory for cheaters was largely due to males' memory for the faces of low status cheaters. However, Barclay and Lalumière (2006) noted that some of the written descriptions did not cleanly fit the categories of "cheater" or "trustworthy", and that any written descriptions are potentially prone to bias because the particular types of "cheating" or "cooperation" might make the actions in one category more salient than in other categories.[1] Using a similar methodology and more controlled descriptions, Barclay and Lalumière found that faces of cheaters were remembered equally as often as faces associated with other information (and interestingly, altruists were remembered best in their second experiment). Baron and Burnstein (2002) found that faces paired with positive and negative character information were remembered equally well after a ten minute delay, and both types of faces were remembered better than faces presented without character information. Perhaps most telling is the case where Chiappe and colleagues (Chiappe et al., 2004) reported that cheaters were recognized more frequently than cooperators after a ten minute delay, but the "degree" of cheating by the cheaters was much higher than the "degree" of cooperation by the cooperators.[2] After correcting this problem in 2005, they found that cheaters and cooperators were remembered equally often (McCulloch & Chiappe, 2005). Most recently, Mehl and Buchner (2008) also found no difference in memory for faces paired with positive, negative, or irrelevant character information.

Two other studies deserve mention and are less prone to the methodological problems associated with using verbal descriptions of cheating and cooperation. Using a hypothetical Prisoner's Dilemma, Oda (1997) found that male defectors were remembered more often than male cooperators after a one-week delay, but there was no effect for female faces, all of which were remembered as often as male defectors. Yamagishi and colleagues (Yamagishi, Tanida, Mashima, Shimoma, & Kanazawa, 2003) found that people tended to remember the faces of people who actually had been defectors in a real Prisoner's Dilemma, but tended to "remember" them even if they had not actually seen them before. This result is more consistent with the idea that there is something about the way defectors look (see also Brown, Palameta, & Moore, 2003) than with special memory for defectors. Together, these results show that the evidence for enhanced recognition of cheaters or defectors is equivocal at best.

[1] One example of a "cheater" from Mealey et al. (1996) is a doctor who sexually abused patients, and an example of a "trustworthy" person is an airport shoe-shiner who noticed and reported a potentially dangerous oil leak in one of the planes.

[2] In the one example in the published paper, the cheating and cooperation is relatively balanced (payment/non-repayment of a debt). However, other items used include unscrupulous real-estate agents who sold houses on former toxic waste dumps to families with children without telling them of the "extraordinarily" high rates of child leukemia, compared to scrupulous agents who did tell the families about the risks (but sold them the houses anyway), or pharmacists who do (versus don't) make extra money selling a narcotic to people without prescriptions, for whom it might fatal. In examples such as these, cheating involves gaining potential profit through gross disregard for others' health, and is surely more salient than not taking the extra profit.

1.2. The current study

The current study tested whether people's social exchange mechanisms are sensitive to the proportion of defectors in a population, such that people remember defectors when cooperators are most common and vice versa. Participants played a trust game with computerized players and saw the faces of 40 "players" who cooperated or defected in this trust game. Defectors comprised 20%, 50%, or 80% of the faces. After a distracter task, participants were asked which faces they remembered and were asked to classify each of them. If people possess a psychology designed for general reputation-tracking (which could apply a domain-general "remember-the-rare-type" strategy for remembering actions) rather than a psychology specifically designed for recalling cheaters, then they should remember whichever type of "player" is rare. If instead, people possess a module for "cheater recognition", then participants should remember defectors best regardless of the number of each type that they see. One must note that these results will not speak to whether people can detect violations of social contracts or recognize instances of defection as proposed by Cosmides and Tooby (e.g. 1992, 2000), because people still need to identify defections and/or cooperation to track reputations. Participants should trust more often when defectors are rare regardless of the specificity of the cognitive mechanism.

2. Methods

2.1. Participants

Participants were students at Cornell University who were recruited via announcements in class and posters. Thirty-six females (mean age 19.2 $\pm$ s.d. 1.9 years) and 24 males (mean age 21.4 $\pm$ s.d. 7.4 years) of mixed ethnicity participated.

2.2. Stimuli

Facial photographs of university-age Caucasian males were obtained from the CVL Face Database (originally described in Solina, Peer, Batageli, Juvan, & Kovac, 2003) and were supplemented with photographs from the Face Research Lab at the University of Aberdeen. From these sets, faces were selected for neutrality of expression and a low number of distinguishing features.

2.3. Trust game

Participants played a trust game in which they were paired with computerized players and could choose whether to trust those players. If a participant did not trust, then the game ended and both players (the participant and the computerized player) earned \$10. If the participant did trust, then the computerized player could choose to cooperate or defect. If the computerized player cooperated, then both players earned \$15, but if it defected, then it would earn \$20 and the participant would earn \$5. Thus, participants could earn more money than they started with if they trusted a cooperator, but would be worse off if they trusted a defector, so there was an advantage to knowing a partner's cooperativeness. Participants had full knowledge that they were playing against computerized players, yet the trust game gave them a vested interest in trusting cooperative computerized players. To better simulate a real trust situation, the computerized players were presented as having "earnings" to provide a rationale as to why one of them might "want" to defect, but the computer players' "earnings" did not affect the participants' outcomes.

2.4. Procedure

Participants saw the faces of 40 computerized players sequentially in a random order and saw whether each one would cooperate or defect in the trust game (cooperation or defection was randomly assigned to each face). Participants had been told beforehand that they would play the trust game with six of the computerized players. Participants rated the attractiveness of each face. There were three experimental conditions based on the proportion of defectors: Defectors comprised 20%, 50%, or 80% of the 40 faces. There were 20 participants in each condition.

After seeing and rating the 40 faces, participants completed demographic questions and personality scales as distracter tasks (data not reported here); these took approximately 10 min to complete. Afterwards, participants were given an unexpected recognition test. They were shown the previous 40 faces along with 40 new ones (in a random sequential order) and were asked to classify each face as previously seen or novel (memory test) and as a cooperator or defector (classification test). Participants received an extra five cents for each face correctly classified. Following this recognition test, participants played the trust game with three cooperators and three defectors, although they did not know how many of each type were there. Participants received the average payoff from the six trust games in addition to the money they earned in the recognition test, and this averaged \$13.58. This research methodology was approved by the Cornell University Committee on Human Subjects. All statistics presented are two-tailed tests.

3. Results

Overall, memory for faces was good (mean = 30.1 out of 40, s.e. = 0.82, median = 32, mode = 33) and false alarms were low (mean = 3.8 out of 40, s.e. = 0.56, median = 2.5, mode = 0), both of which are comparable to the results in Barclay and Lalumière (2006) despite methodological differences. The data were analyzed with the relative proportion of defectors and cooperators remembered as a within-subject factor, and the rarity of defectors as a between-subjects factor. Participants differed in the three conditions in their memory for defectors relative to cooperators ($F_{2,57} = 7.83$, $p = .001$, Fig. 1). Planned contrasts revealed that the 20% defector condition differed significantly from the 50% defector condition ($p = .005$), but the latter did not differ from the 80% defector condition ($p = .39$). Using planned analyses

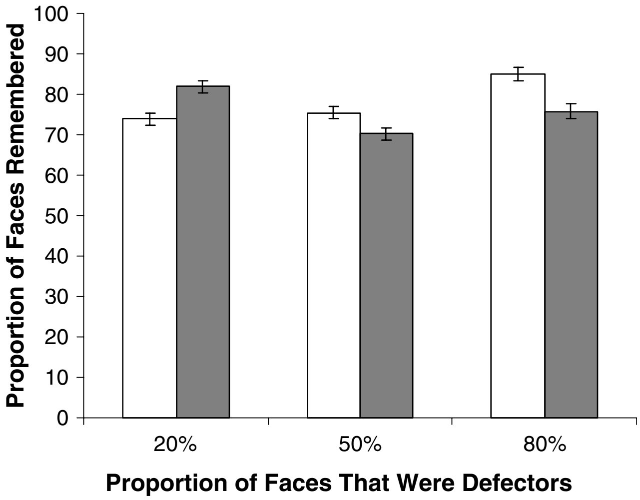

Fig. 1. Average proportion (and standard errors) of cooperators (light bars) and defectors (grey bars) recognized when defectors represented 20%, 50%, and 80% of the faces seen. Within-subject errors bars were calculated according to Cousineau's (2005) method of factoring out individual differences in response.

within each condition, defectors were remembered better than cooperators in the 20% defector condition (paired $t = 2.61$, $p = .017$), but marginally worse than cooperators in the 50% defector condition (paired $t = 1.79$, $p = .09$) and significantly worse than cooperators in the 80% defector condition (paired $t = 2.55$, $p = .02$). There was no main effect of participant sex, nor did it interact with other variables (all Fs < 1), so it will not be discussed further.

When analyzing accuracy in classification, we must consider the probability of correct classifications by chance. If participants randomly classify some faces as defectors in the same proportion in which they saw them earlier, then they will be correct about each type as often as that type appeared. For example, if they originally saw 20% defectors and they later classify a randomly selected 20% of the faces as defectors, then they would make correct classifications for 80% of the cooperators and 20% of the defectors. When participants saw 20% defectors, they later classified defectors more accurately than chance would predict (one sample $t = 4.06$, $p = .001$, Fig. 2) but did not do so for cooperators ($t < 1$). When they saw 50% defectors, cooperators were classified more accurately than chance (one sample $t = 3.37$, $p = .003$) but defectors were only marginally so (one sample $t = 1.89$, $p = .079$). When participants saw 80% defectors, both defectors and cooperators were classified more accurately than chance (one sample ts $= 2.57$ and 3.04, ps $= .019$ and $.007$, respectively). For the previous analysis, the actual a priori proportion of each type may not be the correct benchmark against which to compare participant's success at classifying cooperators and defectors. Alternately, instead of the a priori proportion of each type, we can find the proportion of times that each participant classified any face (previously seen or novel) as a cooperator or defector, and use that a pos-

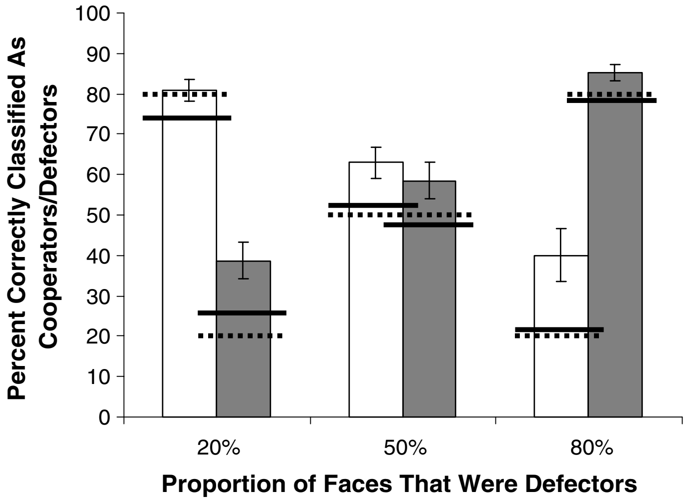

Fig. 2. Average proportion (and standard error) of faces correctly classified as cooperators (white bars) or defectors (grey bars) when defectors represented 20%, 50%, and 80% of the faces seen. The dashed lines represent the likelihood of correctly classifying cooperators and defectors if one guesses in proportion to the proportion of that type among the faces previously seen, whereas the solid lines represent the chance likelihood of correctly classifying cooperators and defectors based on the a posteriori number of faces (old and new) classified by participants as cooperators and defectors.

teriori proportion as the expected chance level of correctly classifying a cooperator or defector for that participant. This controls for participants' beliefs about the proportion of each type. Using this benchmark, participants were more accurate than chance at classifying both cooperators and defectors in the 20% defector condition (one sample ts = 2.73 and 2.71, ps = .013 and .014, respectively, Fig. 2), the 50% defector condition (one sample ts = 2.80 and 2.34, ps = .011 and .031, respectively), and the 80% defector condition (one sample ts = 2.78 and 3.39, ps = .012 and .003, respectively).

The important question is whether participants were more accurate at classifying defectors than cooperators depending on the frequency of each, so we must control for the probability of classifying each type by chance (based on the proportion of times each participant classified any face as a cooperator or defector). By subtracting that expected proportion from the observed proportion of correct classifications, multiplying by 100 and then dividing by the expected proportion according to the formula: [(observed–expected) * 100/(expected)], we produce a percentage measuring how much better than chance participants were at classifying cooperators and defectors. These percentages can be used to compare the accuracy of classifying cooperators and defectors even though those types may have differed in the proportion of faces observed. Participants who saw 20% defectors were slightly but not quite significantly more accurate at classifying defectors than cooperators (average difference in percentage = 57 ± s.e. 32, paired $t = 1.79$, $p = .091$), those who saw 50% defectors were equally accurate at classifying both types (average difference in percentage $= 4 \pm$ s.e. 8, paired $t < 1$, n.s.), while those who saw 80% defectors were significantly more accurate at classifying cooperators (average difference in percentage $= 110 \pm$ s.e. 51, paired $t = 2.16$, $p = .044$). The three conditions differed significantly in the accuracy of cooperators relative to defectors ($F_{2,56} = 5.74$, $p = .005$), and planned contrasts revealed that relative accuracy of cooperators and defectors did not differ between the 20% defectors and 50% defectors conditions ($p = .22$), but the latter did differ from the 80% defector condition ($p = .037$). This suggests that after controlling for the likelihood of classifying any face as a cooperator or defector, people were most accurate at classifying the rare type, but only significantly so for rare cooperators.

If participants had made completely correct trust decisions in the trust game, they would have trusted three out of six times. Participants trusted significantly more often than three times when they saw 20% defectors (average: $3.9 \pm$ s.e. 0.3 times, one sample $t = 3.21$, $p = .005$), non-significantly less than three times when they saw 50% defectors (average: $2.6 \pm$ s.e. 0.26 times, one sample $t = 1.57$, $p = .13$), and marginally less than three times when they saw 80% defectors (average: $2.4 \pm$ s.e. 0.36 times, one sample $t = 1.78$, $p = .091$). Trust levels differed significantly between the three conditions ($F_{2,57} = 7.51$, $p = .001$), and planned contrasts show that the 20% and 50% conditions differed significantly ($p = .004$) whereas the 50% and 80% conditions did not ($p = .56$).

4. Discussion

These results show that enhanced recognition and classification of defectors only occurs when they are rare. When they were common, cooperators were remembered and classified best. When the two types were equally common, they were classified equally well, and cooperators were remembered slightly better than defectors. Although this last finding did not quite reach traditional significance levels, it is similar to the high rates for remembering altruists found in Experiment 2 of Barclay and Lalumière (2006), and is strong evidence against the argument that defectors would be remembered best if both types were equally common. These results suggest that the apparent salience of defectors depends on them being less common than cooperators. This refutes the idea that human cognitive mechanisms for social exchange automatically focus on cheaters, and instead provides evidence for cognitive mechanisms that are slightly more general than that. Such mechanisms might focus attention on others' general social reputation (i.e. cooperation or defection) with the obtained information being remembered by a general memory that notices rarity. This is supported by the fact that although participants were more accurate than chance at classifying both cooperators and defectors, the relative accuracy for each type also depended somewhat upon the rarity of each type. Although this study did not directly test whether people do pay particular attention to information about others' reputation compared with other types of information, one would predict a bias towards tracking any information that could influence one's future well-being, such as others' past actions. Indeed, other research shows that people do attend to the cooperation of others and base decisions upon that (e.g. Barclay, 2004; Wedekind & Milinski, 2000) and that gossip about social topics (including others' reputation) is a main focus of people's conversation (Dunbar, 2004).

One could argue that general reputation-tracking mechanisms could be more cognitively demanding and error-prone than cheater-specific mechanisms in that they rely on knowledge of the proportion of defectors in a population or group to track one type more than the other. However, if one has not observed or interacted with enough people to know the composition of the population, then there is very little cognitive load and chance of error simply because of the low number of interactions one has had or seen. Furthermore, it is useful to know the probability of encountering cooperation regardless of whether one focuses on cooperators or defectors, so one must track the proportions of each regardless of the specificity of one's social exchange mechanisms. The current data suggest that participants' trust is indeed responsive to the likelihood of encountering cooperation. Memory for rare types can emerge after sampling a population if distinctiveness effects result from memory retrieval instead of the encoding process (McDaniel & Geraci, 2006).

A potential weakness in this study is that participants saw faces of computerized participants rather than actual cheaters with whom participants would interact, such that any cheater-specific adaptations may not have been activated. However, other studies used similar methodology and claimed to find special memory for cheaters (Chiappe et al., 2004; Mealey et al., 1996; Oda, 1997). If an adaptation for automatically focusing on cheaters did exist and could be triggered by those studies or by written scenarios in the Wason Selection Task (as argued by Cosmides & Tooby, 1992, 2000), then it should also be triggered by the cooperative task in the present study. Such a putative adaptation should be activated (albeit, not as strongly) even when participants are not interacting with real cheaters, just as men's sexual arousal can be triggered by pornographic images that are not real women, reputation-seeking mechanisms can be activated by stylized eyespots on a computer (Haley & Fessler, 2005), and kin-recognition systems can be activated towards non-kin by subtle cues of facial resemblance (DeBruine, 2005). In this study, defectors were recalled less well than cooperators when both were equally frequent: If it exists, the putative module was not even partially activated, suggesting that it does not in fact exist.

This study may help broaden the discussion surrounding cognitive adaptations for social exchange, bringing it away from cheaters in particular and towards reputation in general. People still need to recognize when social contracts have been kept and broken (Cosmides & Tooby, 1992, 2000, 2005), but preferential interaction with cooperators could sometimes be better accomplished by focusing on cooperators instead of defectors. An overemphasis on cheating and defection ignores the costs and benefits of investing cognitive resources in detecting particular outcomes and in memory for particular types of people. As such, it would be useful to consider the benefits people may get from focusing on other aspects of cooperation such as the altruistic tendencies of other people or the particular gains and losses of cooperation in different situations.

Acknowledgements

I thank Peter Peer for providing access to the CVL Face Database and Lisa DeBruine for additional faces, Brian Mayer for helping collect data, H. Kern Reeve, Martin Daly, Margo Wilson, and Daniel Krupp for valuable discussions, and three anonymous reviewers for comments. This project was funded by the Department of Neurobiology & Behavior at Cornell University.

References

Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science, 211*, 1390–1396.

Barclay, P. (2004). Trustworthiness and competitive altruism can also solve the tragedy of the commons. *Evolution & Human Behavior, 25*, 209–220.

Barclay, P., & Lalumière, M. L. (2006). Do people differentially remember cheaters? *Human Nature, 17*, 98–113.

Barclay, P., & Willer, R. (2007). Partner choice creates competitive altruism in humans. *Proceedings: Biological Sciences, 274*, 749–753.

Baron, A.S., & Burnstein, E.U. (2002). Are humans equipped with a specific cheater-detection module or a more general person-impression module: Evidence from a face recognition experiment. *Poster presented at the 14th Annual Meeting of the Human Behavior and Evolution Society,* Rutgers University, New Jersey (June, 2002).

Barrett, L., Dunbar, R., & Lycett, J. (2002). *Human evolutionary psychology*. Princeton, NJ: Princeton University Press.

Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocation, and cooperation. *American Economic Review, 90*, 166–193.

Brown, W. M., & Moore, C. (2000). Is prospective altruist-detection an evolved solution to the adaptive problem of subtle cheating in cooperative ventures? Supportive evidence using the Wason selection task. *Evolution and Human Behavior, 21*, 25–37.

Brown, W. M., Palameta, B., & Moore, C. (2003). Are there nonverbal cues to commitment? An exploratory study using the zero-acquaintance video presentation paradigm. *Evolutionary Psychology, 1*, 42–69.

Chiappe, D., Brown, D., Dow, B., Koontz, J., Rodriguez, M., & McCulloch, K. (2004). Cheaters are looked at longer and remembered better than cooperators in social exchange situations. *Evolutionary Psychology, 2*, 108–120.

Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The Adapted Mind* (pp. 163–228). New York, NY: Oxford University Press.

Cosmides, L., & Tooby, J. (2000). The cognitive neuroscience of social reasoning. In M. S. Gazzaniga (Ed.), *The New Cognitive Neurosciences* (2nd ed., pp. 1259–1270). Cambridge, MA: MIT Press.

Cosmides, L., & Tooby, J. (2005). Neurocognitive adaptations for social exchange. In D. Buss (Ed.), *The Handbook of Evolutionary Psychology* (pp. 584–627). Hoboken, NJ: John Wiley & Sons.

Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology, 1*, 42–45.

DeBruine, L. M. (2005). Trustworthy but not lustworthy: Context-specific effects of facial resemblance. *Proceedings: Biological Sciences, 272*, 919–922.

Dunbar, R. I. M. (2004). Gossip in evolutionary perspective. *Review of General Psychology, 8*(2), 100–110.

Fehr, E., & Fischbacher, U. (2004). Social norms and human cooperation. *Trends in Cognitive Sciences, 8*, 185–190.

Graesser, A. C., Woll, S. B., Kowalski, D. J., & Smith, D. A. (1980). Memory for typical and atypical actions in scripted activities. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 503–515.

Haley, K. J., & Fessler, D. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution & Human Behavior, 26*, 245–256.

Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., et al. (2005). Economic man in cross-cultural perspective: Behavioral experiments in 15 small-scale societies. *Behavioral & Brain Sciences, 28*, 795–855.

Hunt, R. R. (1995). The subtlety of distinctiveness: What van Restorff actually did. *Psychonomic Bulletin and Review, 2*(1), 105–112.

Hunt, R. R. (2006). The concept of distinctiveness in memory research. In R. R. Hunt & J. B. Worthen (Eds.), *Distinctiveness and memory* (pp. 3–26). New York, NY: Oxford University Press.

Komorita, S. S., & Parks, C. D. (1995). Interpersonal relations: Mixed-motive interaction. *Annual Review of Psychology, 46*, 183–207.

McCulloch, K., & Chiappe, D. (2005). Memory for cheaters and cooperators in social contract situations. *Poster presented at the 17th annual meeting of the human behavior and evolution society,* Austin, Texas (June, 2005).

McDaniel, M. A., & Geraci, L. (2006). Encoding and retrieval processes in distinctiveness effects: Towards an integrative framework. In R. R. Hunt & J. B. Worthen (Eds.), *Distinctiveness and memory* (pp. 65–88). New York, NY: Oxford University Press.

Mealey, L., Daood, C., & Krage, M. (1996). Enhanced memory for faces of cheaters. *Evolution and Human Behavior, 17*, 119–128.

Mehl, B., & Buchner, A. (2008). No enhanced memory for faces of cheaters. *Evolution and Human Behavior, 29*(1), 35–41.

Nelson, W. R. Jr., (2002). Equity or intention: It is the thought that counts. *Journal of Economic Behavior and Organization, 48*, 423–430.

Oda, R. (1997). Biased face recognition in the Prisoner's Dilemma game. *Evolution and Human Behavior, 18*, 309–315.

Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. Chicago, IL: University of Chicago Press.

Roberts, G., & Renwick, J. S. (2003). The development of cooperative relationships: An experiment. *Proceedings: Biological sciences, 270*, 2279–2283.

Solina, F., Peer, P. Batageli, B., Juvan, S., & Kovac, J. (2003). Color-based face detection in the 15 seconds of fame art installation. In Mirage 2003, conference on computer vision/computer graphics collaboration for model-based imaging, rendering, image analysis and graphical special effects, March 10–11 2003, INRIA Rocquencourt, France, Wilfried Philips, Rocquencourt, INRIA, 2003, pp. 38–47.

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology, 46*, 35–57.

Wedekind, C., & Milinski, M. (2000). Cooperation through image scoring in humans. *Science, 288*, 850–852.

Yamagishi, T., Tanida, S., Mashima, R., Shimoma, E., & Kanazawa, S. (2003). You can judge a book by its cover: Evidence that cheaters may look different from cooperators. *Evolution & Human Behavior, 24*(4), 290–301.