

EXPLORING A FUNCTIONALIST MODEL OF MORAL ELEVATION

A thesis submitted for the degree of Doctor of Philosophy

By

Amy Monroe

**Department of Life Sciences, Division of Psychology
Brunel University London**

Acknowledgements

I would like to thank my supervisor, Michael Price, for his support, feedback, and endless patience in the face of my unusual working habits - I certainly would not have completed this work without his understanding and help. I would also like to thank my loving parents for believing in me, and in particular my mother for lending me her proof-reading skills.

My thanks to the wonderfully encouraging Gail Corrin, and University College Isle of Man for facilitating data collection. Finally, and most gratefully, to my excellent friend Martin Kerr, for always being on hand with every conceivable kind of help.

Abstract

Moral elevation is a positive emotional response to the benevolent behaviour of third-parties, and it is known to engender a motivation to behave benevolently towards others in a remarkably general sense. This ‘contagious’ benevolence is of obvious social value, but psychological research into moral elevation has been underway for less than two decades; consequently, relatively little is known at present about the nuances of the emotion’s elicitors and effects. This thesis addresses some of the gaps in our understanding of moral elevation, by investigating the possibility that moral elevation may be a naturally selected cognitive adaptation. Specifically, this thesis attempts to uncover whether moral elevation’s observable characteristics present evidence of special design; that is, evidence that the emotion operates in such a way as to suggest that it reliably and efficiently performs a particular adaptive function.

Two theories about moral elevation’s possible adaptive function are described: the relationship-building hypothesis, which suggests that elevation serves to aid us in recognizing and capitalizing on opportunities to build high-quality social relationships; and the reputation-management hypothesis, which suggests that elevation helps us to ensure that we keep abreast of the competition to attract and retain social relationship-partners. Although these two hypotheses are similar, and make largely identical predictions about the characteristics moral elevation should be expected to exhibit, *prima facie* reasons for preferring the reputation-management hypothesis are explained.

Tests of predictions derived from these theories of function are reported, and the results of these tests offer some support for the idea that moral elevation could be a functionally specialized cognitive adaptation. Moral elevation appears to be elicited by behaviours which indicate that a third-party has attractive social-partner qualities – namely, a tendency to behave benevolently, a tendency to value the welfare of others, and the potential to generate positive externalities. Elevation also seems to produce behaviours which would increase the social attractiveness of the elevated individual – namely, an increased willingness to value the welfare of others, and strengthened motivations to engage in a variety of pro-social behaviours.

The results presented here suggest that moral elevation has a reliable pattern of activation, potentially indicating that it is a functionally integrated part of our social-behaviour-regulating cognitive architecture. Limitations of the thesis are discussed, and potentially fruitful avenues for further research are considered.

Contents

Chapter 1: Introduction	1
Cognitive Adaptations for Regulating Benevolent Behaviour	2
Moral Elevation	5
Identifying Adaptations	7
Emotions as Cognitive Mechanisms	9
Moral Elevation: Theories About Function	13
The relationship-building hypothesis	14
The reputation-management hypothesis	16
Generating predictions about form	18
Chapter 2: Moral Elevation’s Elicitors	21
Literature Review	22
Morality vs. Benevolence	23
Benefit Delivery: Ability vs. Willingness	26
Incentives to Trade-off Welfare	29
Indirect Benefit Delivery	30
Chapter 3: Empirical Tests of Predictions 1 and 2	32
Study 1: Is Benefit Conferral Necessary to Elicit Moral Elevation?	33
Study 2: Is Moral Elevation Sensitive to Proportionate Cost Incurred By The Actor?	40
Conclusion	45
Chapter 4: Empirical Tests of Predictions 3 and 4	47
Studies 3A and 3B: Is Moral Elevation Sensitive to The Presence of Incentives for The Benefactor’s Behaviour?	48
Study 4: Is Moral Elevation Sensitive to Socially-Attractive Characteristics of The Actor?	62
Conclusion and General Discussion	73
Chapter 5: Moral Elevation’s Effects	76
Recalibrational Emotions	77
Up-Regulating Extrinsic Welfare Trade-Off Ratios.....	80
Interaction With Other Attractive Partner Qualities	81
Motivational Priorities	81

Chapter 6: Empirical Tests of Predictions 5 and 6	84
Study 5: Does Moral Elevation Increase Extrinsic WTR, And If So, Is This Effect Attenuated by High Partner-Value?	85
Study 6: Do Elevated Individuals Show Particular Motivational Preferences?	99
Omnibus Analysis: Sex and Religiosity Differences Across Studies	105
Chapter 7: Conclusions and Future Research Directions	107
Evidence That Moral Elevation’s Inputs Conform to the Expectations of a Functionalist Analysis	108
Evidence That Moral Elevation’s Outputs Conform to the Expectations of a Functionalist Analysis	110
Evidence of Sex and Religiosity Differences in Moral Elevation	112
Conclusion	114
Cross-cultural comparisons	114
Individual differences	115
Practical and theoretical utility	116
Appendix A: Study 1 Materials	126
Appendix B: Study 3A Materials	132
Appendix C: Study 3 B Materials	138
Appendix D: Study 4 Materials	143
Appendix E: Study 5 Materials	144
Appendix F: Study 6 Materials	148

Chapter 1: Introduction

Abstract

This chapter introduces the subject of interest for the thesis – a social emotion called moral elevation. The first section gives a brief summary of evolutionary explanations for human prosocial behaviour, and locates the thesis' topic within a broader framework of current academic enquiry. The second section summarizes the extant moral elevation literature and identifies a gap in the literature concerning whether the emotion may have been sculpted by natural selection. A précis of the criteria for labeling a trait an adaptation is offered, and a framework for understanding emotions as evolved cognitive mechanisms is outlined. Two theories about the possible adaptive function performed by moral elevation are explained; the relationship-building hypothesis, and the reputation-management hypothesis. *Prima facie* reasons for preferring the latter over the former are laid out, and the chapter ends with a summary of predictions about elevation's form, which are to be expanded on in subsequent chapters.

Cognitive Adaptations for Regulating Benevolent Behaviour

Humans are a uniquely benevolent species; we have an unparalleled propensity to deliberately increase the welfare of others, even at a cost to ourselves (Høgh-Olesen, 2010; Katz, 2000; Kurzban, Burton-Chellew, & West, 2015; Wilson, 2012). Whether it be spending time or energy, sharing resources, passing up opportunities, or compromising our safety, we routinely incur various costs with the intention of benefitting other individuals.

This wide-spread benevolence was initially viewed by evolutionary theorists as a serious challenge to the idea of natural selection (Fletcher & Doebeli, 2009). No trait can reach fixation in a population if it diminishes the fitness of the individuals who express it, and absent any other qualifiers, benevolent behaviour is definitionally fitness retarding. But over the past 60 years researchers have described a variety of circumstances in which the direct costs of ceding welfare to others can be recouped by the actor to generate a net fitness gain (West, Griffin, & Gardner, 2007).

The first proposal of a circumstance in which benevolent behaviour could be fitness-promoting for the actor came in the form of Hamilton's (1964) kin selection theory. Hamilton highlighted the idea of inclusive fitness, the central insight of which is that other individuals share copies of some of our genes by virtue of common descent. If a gene, or more likely a set of genes, codes for behaviours which promote the reproductive interests of other individuals who carry them, they can positively influence their relative representation in future generations even if they result in an overall relative reproductive disadvantage for the individual they express themselves in. So long as the fitness advantage created for other copies of these benevolent-behaviour-producing genes is larger than the direct fitness cost they create for the individual carrying them, an inclusive fitness gain is generated, and they can increase in population frequency across generations.

Inclusive fitness may be conceptually necessary for explaining the evolution of benevolent behaviours which generate a life-time direct fitness cost for the actor, but as Trivers (1971) pointed out, not all direct fitness costs last a life-time. Sequential acts of benevolence between two individuals can, over time, produce gains in trade which will increase the reproductive success of both parties. So long as the odds of reciprocation are high, and the costs incurred by the benefactor in any benevolent instance are smaller than the benefits produced for the recipient, directly-reciprocal

welfare-ceding behaviour can generate a life-time direct fitness gain for the individuals who engages in it (Trivers, 1971; Axelrod & Hamilton, 1981).

Parties of three or more individuals can further increase the scope for generating gains in trade, if they engage in indirectly-reciprocal benevolence (Alexander, 1987; Axelrod & Hamilton, 1981; Boyd, Gintis, & Bowles, 2010; Nowak & Sigmund, 2005; Trivers, 1971). A network of indirect-reciprocity operates so that each member of the network will give assistance to any other network member, so long as that member has a prior history of helping others herself (Alexander, 1987; Panchanathan & Boyd, 2004). Restricting mutual-benevolence to direct, two-party exchanges carries with it the risk of failures in reciprocity resulting from unpredictable natural events; in any non-simultaneous dyadic exchange, there is a risk that your exchange partner may be unable to reciprocate when it would benefit you most (they may be sick, have fallen on hard times, or be out of town, etc.). This risk is distributed, and so minimized, in systems of indirect-reciprocity – even if the individual you originally helped is unable to help you back, another member of the network will likely be able to offer assistance when needed (Wedekind & Braithwaite, 2002). Again, so long as each member of the network restricts their benevolence to situations in which the costs they incur are outweighed by the benefits they produce for others, the overall effect is a net direct fitness gain for each.

Direct fitness gains from benevolence need not come just from reciprocity. Displays of benevolent behaviour can be used as costly signals of an individual's skills, qualities, and resources, which promote the individual's fitness interests by increasing the chances that they will be chosen by others for valuable reproductive or social relationships (Fehr & Fischbacher, 2004; Noë & Hammerstein, 1995; Roberts, 1998). The resulting reproductive advantages can outweigh the costs of benevolent behaviour even in the absence of reciprocity.

Strategically ceding contested resources can also act as a fitness promoting conflict resolution strategy. When another individual would be willing to fight for ownership of some resource (be it your food, your mate, or your social position), and when the value of that resource to you is lower than the cost the other would likely be able to inflict in combat, the least fitness-damaging behavioural option available to you will be to willingly cede what the other individual demands (Hammerstein & Parker, 1982).

So long as benevolent behaviour generates greater benefits for the recipient than the costs the actor incurs in performing it, and so long as its expression is contingent on the probability that one of the gains outlined above will eventuate, it can be a fitness promoting behavioural tendency, capable of spreading through a population by natural selection (Delton, 2010).

A slew of results in the empirical literature support the prediction that human benevolence should be dependent on these factors pertaining. Across a variety of laboratory and naturalistic settings, people have proved markedly more willing to benefit others when fitness incentives exist for them to do so (or, more accurately, when there would have been fitness incentives for our ancestors in the environment of evolutionary adaptedness) (e.g. Barclay, 2010; Barclay & Willer, 2007; Chaudhuri, 2011; Curry, Roberts, & Dunbar, 2013; Fehr & Gächter, 1999; Fischbacher, Gächter, & Fehr, 2001; Gächter, 2007; Herrmann, Thoni, & Gächter, 2008; Price, 2003; Raihani & Smith, 2015). Humans do indeed seem to adaptively regulate their engagement in benevolent behaviour so as to selectively exploit those circumstances in which it would, ancestrally and on average, have produced some kind of fitness gain for the actor. Despite being prolific, then, our tendency to cede welfare to others is far from random or indiscriminate. Indeed, it has become such a well-displayed trait precisely because our species has grown so adept at being discerning about how much welfare we cede to whom and when.

This careful regulation is not achieved through a process of conscious deliberation. In our intensely social lives, we face a host of situations every day where we have to choose between benefitting ourselves or benefitting others. Do we forgo a second cup of coffee so there will be enough milk left for someone else to have one? Do we bring in wood for the fire, or wait for someone else to get cold and do it first? Do we humour the office bore at the Christmas party by listening to his thoughts on wind-farms? For how long – five minutes? Ten? For the most part we are unaware of these decisions even being made, but navigating all of these choices adaptively (i.e., ceding welfare only when it is likely to be fitness promoting to do so) is arguably our evolutionary wheelhouse, the skill that allows us to occupy the cooperative niche that we do (Gintis, 2003; Høgh-Olesen, 2010; Krebs, 2011; Wilson, 2012).

Our ability to adaptively manage when and to what extent we cede welfare to (and impose costs on) others rests on a specialized cognitive decision-making

architecture. Natural selection has equipped us with an intricate set of mental mechanisms, which function in concert, allowing us to increase and decrease our inclination to behave benevolently as situations (and therefore incentive structures) change (Cosmides & Tooby, 1994; Krebs, 2008; Barkow, Cosmides, & Tooby, 1992). Turning a fitness profit from intense sociality carries with it a complex set of information processing problems, and we are only able to successfully circumvent these problems because we have evolved a complex set of cognitive devices, each one specialized to overcome a specific difficulty.

Evolutionary psychologists have recently begun to map the computational properties of this decision-making architecture, bringing to light the variety of mental mechanisms involved in regulating our willingness to cede welfare to others. Mechanisms specialized for gathering, storing, and summarizing specific types of information, monitoring the environment for changes in relevant mutable variables, and regulating motivation and behaviour; all coordinating to solve different aspects of the same overall problem - living in an ecology where our welfare and the welfare of others often overlap or conflict (Cosmides & Tooby, 1992; Delton, Cosmides, Guemo, Robertson, & Tooby, 2012; Delton & Robertson, 2012; Lieberman, Tooby, & Cosmides, 2007; Lim, 2012; Sell et al., 2009; Sell, Tooby, & Cosmides, 2008; Tooby & Cosmides, 2008; Tooby, Cosmides, Sell, Lieberman, & Sznycer, 2008).

The central purpose of this thesis is to explore the possibility that a little-studied emotion called moral elevation may be another member of this set of adaptations. Although it was only recently described as an object of research, uptake in interest in the emotion has been relatively quick, particularly in the field of positive psychology, because it seems to increase people's willingness to engage in benevolent behaviour.

Moral Elevation

Moral elevation is the warm, uplifting emotional sensation that we sometimes experience when we witness or hear stories about others' benevolent actions (Haidt, 2003a). It can be accompanied by a feeling of poignancy (Piper, Saslow, & Saturn, 2015) and a sensation of warmth in the chest (Algoe & Haidt, 2009), and it invariably results in a broad, pleasant feeling, of wanting to behave more benevolently towards other people ourselves (Haidt, 2000).

The first academic description of moral elevation was given by Jonathan Haidt, who coined the name as a way to highlight a proposed symmetry with socio-moral disgust (Haidt, 2000). He suggests that, whereas seeing others behave in socially undesirable ways can lead us to judge those others, and humanity more generally, as base and aversive, seeing others behave in a socially positive manner can lead us to judge those others, and humanity more generally, as pure and attractive. Placing moral disgust and moral elevation at opposite ends of a 'vertical continuum', Haidt suggests that disgust is provoked by displays of humanity's worse, lower nature, and elevation is provoked by displays of humanity's better, higher nature (Haidt, 2003a, 2003b, 2010).

Although moral elevation is similar in many respects to gratitude, it has two substantive differences. The first is that, unlike gratitude, elevation is sensitive to 'disinterested elicitors' (Haidt, 2003b) - others' behaviours which do not directly affect the self. Whereas gratitude is provoked when the actions of another benefit oneself, moral elevation is experienced in response to "the sight of a stranger helping another stranger" (Haidt, 2010, p.86).

The second differentiator is the scope of the motivational outputs that the two emotions produce. Although both gratitude and elevation induce a desire to benefit the party whose actions evoked the emotion, elevation also creates a desire to benefit others more broadly (Haidt, 2003a; Algoe & Haidt, 2009). While the benevolence that gratitude motivates is targeted specifically at the individual to whom one is grateful, the motivation to behave benevolently that elevation produces appears to be relatively indiscriminate. Extending his analogy with disgust, Haidt describes this 'prosocial action tendency' as a kind of 'positive contamination' (Haidt, 2003b).

Experiences of moral elevation have been found to increase rates of volunteering (Cox, 2010), organizational citizenship behaviour (Vianello, Galliani, & Haidt, 2010), helping (Schnall, Roper, & Fessler, 2010), mentoring (Thomson, Nakamura, Siegel, & Csikszentmihalyi, 2014), and charitable donation (Van de Vyver & Abrams, 2015). Proneness to experiences of moral elevation has also been found to correlate with self-reported prosocial behaviour, showing incremental predictive validity over the five-factor personality model (Landis et al., 2009). There is even evidence that elevation can motivate us to behave benevolently toward people to whom we may previously have been hostile (Freeman, Aquino, & McFerran, 2009),

and that it can reduce expressions of prejudice, both racial (Freeman et al., 2009) and homophobic (Lai, Haidt, & Nosek, 2014).

Despite burgeoning interest, research into moral elevation is still in its infancy. As yet there have been no attempts to offer an evolutionarily grounded analysis of the emotion, and no attempts to formally explore the possibility that elevation may be an adaptive feature of our species-typical cognitive architecture. This thesis intends to address that research gap. A better understanding of what role, if any, this emotion plays in our evolved psychological repertoire would not only further our understanding of how humans are adapted to fill the intensely social niche that we do, it could also facilitate the cultivation of positive social outcomes. A nuanced understanding of how and why we experience 'prosocial contagion' is a necessary precondition for realizing the potential that this phenomenon holds as a tool for social intervention.

The remainder of this chapter will lay out how we can go about investigating whether moral elevation may be an adaptation. The next section will review the criteria which must be met in order for any trait to be labeled an adaptation. The following section outlines the terms in which we can understand and discuss emotions as information processing adaptations. The final section of this chapter reviews two hypotheses about the adaptive function which moral elevation may have been selected to perform.

Identifying Adaptations

Confidence that a behavioural tendency is the output of a specialized cognitive adaptation requires particular standards of proof. Any given phenotypic characteristic (any distinct trait, or 'thing' we can point to in an organism) can be explained by one of three mutually exclusive and collectively exhaustive factors (Buss, Haselton, Shackelford, Bleske, & Wakefield, 1998):

- i) it might be an adaptation, naturally selected because it had a positive effect on the reproductive success of the organism's ancestors

- ii) it might be a byproduct of an adaptation – a characteristic which arises because it is a necessary conjunct of some other naturally selected feature, or

iii) it might be random noise – a characteristic arising from chance mutations in the organism's genotype

To give an example, having rough, ridged skin on the tips of our fingers is an adaptation; it increases our ability to grip and manipulate objects (Adams et al., 2013). The tendency for residues to stick more easily to our fingertips than to other parts of our skin is a byproduct of this adaptation for grip; it is an unavoidable but non-functional consequence of having a grip-enhancing texture on our fingers. Our fingerprints (the precise pattern of the rough ridges on our fingertips) are noise; they vary person to person for no reason other than inconsequential differences between our genotypes.

We can broadly discriminate between adaptations and byproducts on the one hand and random noise on the other by looking for evidence of species-typicality. The majority of adaptations increase in representation in a population, as they outcompete less effective alternative designs, eventually reaching fixation (although sex-selected adaptations and adaptations arising as the result of negative frequency-dependent selection are exceptions (Ayala & Campbell, 1974). Traits which have not been naturally selected (or which are not necessarily linked to those which have been) – those we would characterize as random noise – do not reach fixation. With no impact on relative reproductive success, traits arising from genetic noise simply drift in population frequency by chance (Masel, 2011).

We gain confidence that a psychological trait is part of our reliably developing, species-typical cognitive architecture when we find it occurring across a variety of different cultures. Traits needn't occur in all cultures in order to be thought species-typical - many adaptations are facultative and will develop only when certain environmental contingencies pertain – but finding that a behavioural phenomenon occurs in a range of cultures adds significant weight to the supposition that it is a species-typical characteristic.

The question of moral elevation's culture-independence has yet to be properly addressed in the literature, and the overwhelming majority of empirical studies have relied on WEIRD participants (Western, educated, industrialized, rich, and democratic) (Henrich, Heine, & Norenzayan, 2010). Two small interview-based studies conducted in India and Japan did find that subjects described experiencing emotional reactions with the hallmarks of moral elevation when asked to recall

seeing another person perform a good deed for a third party (Haidt, 2003a). But with their very small sample sizes (8 and 15 respectively) these studies give only a very preliminary hint at moral elevation's species-typicality.

Although evidence of cross-cultural consistency weighs against a trait being the product of random noise, it does not help us to distinguish between adaptations and byproducts. To make the case that moral elevation is an adaptation, we would need to find signs of 'special design' (Williams, 1966). The important difference between adaptations on the one hand and byproducts and noise on the other is that adaptations are functionally organized. They are an arrangement of parts which operate in such a way as to obfuscate an identifiable, reliably recurring, ancestrally relevant obstacle to relative reproductive success. That is to say, adaptations are those features of an organism's phenotype which are so improbably well-suited to solving an adaptive problem faced by the organism's ancestors as to make it unlikely that they would have arisen by chance alone (Buss et al., 1998; Tooby & Cosmides, 1990a). This kind of complex functional organization is known as 'special design', and its hallmarks are efficiency, complexity, economy, reliability, and precision (Williams, 1966). Calling a trait an adaptation requires us to demonstrate that there is a high degree of fit between its form (how it operates) and the function it is proposed to have been selected for.

Emotions as Cognitive Mechanisms

Cognitive adaptations are information processing devices, and so describing their form involves detailing their computational properties – the set of if-then rules that govern their operation. A cognitive mechanism is a structure, instantiated in the brain, which is sensitive to specific forms of input (Barkow et al., 1992; Confer et al., 2010). When an individual's environment presents information which meets the relevant criteria, the mechanism generates a particular set of outputs (an 'output' here being any state-change, whether it be in physiology, affect, cognition, motor-movements, endocrinology, etc.). A single cognitive mechanism may produce a range of conditional outputs which depend on variations in the informational input. These input-output relationships are what constitute the form of the cognitive mechanism.

Describing emotions as the outputs of cognitive mechanisms may seem counterintuitive because it stands at odds with the folk distinction between cognition

and affect ('thinking' versus 'feeling'). But from an evolutionary perspective the brain is an evolved information processing organ, and any and every mental operation must be describable in the computational language of cognitive psychology (Barrett, 2005; Cosmides & Tooby, 1994). This approach does not entail reducing emotional experiences to anything 'deliberative', 'rational', or 'cold' – rather, it allows us to detail the underlying information processing rules which give our emotions their consistently patterned structure.

From an evolutionary perspective, the function and computational structure of emotions is best thought of as a form of overarching mechanism coordination (Al-Shawaf, Conroy-Beam, Asao, & Buss, 2015; Cosmides & Tooby, 2000; Haselton & Ketelaar, 2006; Tooby & Cosmides, 1990b). Our minds house a huge number of specialized cognitive devices, each selected to execute a specific adaptive function. Coordinating the operations of these many separate adaptations is crucial to avoiding two types of problem: conflicting effects, and compound adaptive challenges (Cosmides & Tooby, 2000, 2013).

The problem of conflicting effects arises from the sheer volume and variety of cognitive mechanisms we possess. Each mechanism functions to circumvent a particular adaptive challenge, but we do not face all of our adaptive challenges at once and the solutions to different challenges can be mutually inconsistent (e.g., sleep and predator avoidance, eating and avoiding contamination, social bonding and avoiding social exploitation). If our minds were structured in such a way that any cognitive mechanism could operate unrestricted, regardless of which other mechanisms might be operating, serious fitness-retarding consequences would quickly ensue, with the outputs of one mechanism cancelling out the effects of the next.

Equally importantly, adaptive challenges often present themselves in reliably-patterned clusters, creating compound challenges which necessitate coordination between particular sub-groups of cognitive mechanisms. Our ancestors routinely faced richly structured, fitness-relevant situations which could only be effectively navigated by the employment of a sustained and coordinated behavioural strategy. Situations like being stalked by a predator, meeting an attractive potential mate, finding our offspring in distress, having our social status challenged, and so on, require us to execute a range of specific tasks, some simultaneous, some sequential. The execution of each task is controlled by a particular specialized

cognitive device, and these devices need to work in concert in order to produce a behaviourally coherent response to a complex situation.

Emotions are cognitive programs which evolved to fill this role of superordinate mechanism orchestration (Tooby & Cosmides, 1990b). Each of our emotions responds to a specific, evolutionarily recurrent situation type. The inputs which trigger emotions are cues which reliably indicated, ancestrally, that a situation of the type they evolved to deal with currently pertains. When triggered by such cues, an emotion will mobilize a specific subset of our cognitive architecture's other programs, prompting them to operate in a particular configuration (up-regulating some, down-regulating others, altering the operating parameters of others still, etc.). Emotions initiate a pattern of activation across a varied range of different cognitive mechanisms, which may include:

Those 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; behavioural decision rules; motor systems; communication processes; energy level and effort allocation; affective coloration of events and stimuli; and the 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. (Tooby & Cosmides, 2008, p.118).

The psychological configuration that any particular emotion initiates is the configuration which led, on average, to the most fitness promoting outcomes for our ancestors when they faced situations of the given type. For example, when we experience fear in a situation where we might be being stalked by a predator: our auditory attention threshold shifts, making us more readily aware of noises which might indicate the presence of a threat; the threshold at which we respond to stimuli as though they do constitute a threat is lowered; our goal focus changes, shifting the value we place on different potential outcomes (being somewhere safe or brightly lit suddenly becomes far more attractive and occupying a goal than finding food); and so forth. (Cosmides & Tooby, 2000). Similarly, when we face situations of potential toxic contamination, disgust initiates a pattern of cognitive activity which helps us to coordinate our behaviours in contamination-avoiding ways (Cosmides & Tooby, 2000; Kelly, 2011). Gratitude, guilt, and anger allow us to mobilize adaptive

behavioural routines in situations where we have a chance to consolidate a beneficial social relationship, risk undervaluing an important social partner, or risk being under-valued by another (respectively) (Cosmides & Tooby, 2013; Lim, 2012; Sell, 2005).

In many respects, this adaptationist approach to describing emotions does not stand at odds with or radically depart from other historically important theories of emotion. It is not the first theory to acknowledge that emotions function to help us navigate adaptively important and specific situations (e.g. Ekman, 1999; Oatley & Johnson-Laird, 2014), and neither is it the first to recognise that emotions produce a 'cascade of changes' (Ekman & Cordaro, 2011, p.306), altering a range of other mental functions so as to achieve a particular set of situation-relevant outcomes (e.g., Keltner & Cordaro, 2015; Lazarus, 1991; Oatley & Johnson-Laird, 2014).

Its emphasis on the function of *coordinating* these changes, however, means that the computational approach avoids arguments over the causal and taxonomic primacy of the evaluative, physiological and phenomenological, or expressive and behavioural aspects of emotion experiences, which have predominated in emotion theory over the past fifty years. There are different ways of cutting up the emotion literature (Gendron & Feldman Barrett, 2009; Scarantino & de Sousa, 2018), but theorists have largely agreed that the same set of central elements characterize emotional experiences, while debating which of these elements constitute the necessary and sufficient conditions for an experience to be properly classified as emotional. Some have placed their theoretical emphasis on the feelings – the physiological and phenomenological aspects – of emotional experiences (James, 1884; Schachter & Singer, 1962), others on the evaluative (or appraisal) processes surrounding the triggering of an emotion (Lazarus, 1991; Oatley & Johnson-Laird, 2014), and others still on the motivations and behavioural consequents of emotional experiences (Ekman & Friesen, 1971; Keltner & Cordaro, 2015).

The computational approach maintains all of the same central elements as these traditional theories, but subverts the need to argue about causal primacy by parsing the emotion itself as the superordinate cognitive device which mobilizes other parts of our cognitive architecture to operate in a specific configuration. Each of these superordinate devices is equipped with a 'situation-detecting' component, so that situations which exhibit the requisite cues (likely to be a combination of exogenous and endogenous) will initiate the relevant emotion mechanism, in turn

prompting any necessary changes in physiology, appraisal, cognition, motivation, behaviour, communication, etc. When any superordinate cognitive device of this kind is given to be an emotion, the question over the temporal primacy of the effects it produces is removed (they should be interpreted as functionally simultaneous), and the taxonomic necessity of any particular emotion characteristic, or effect, becomes unimportant, as emotion devices selected to guide us through different types of situations should be expected to produce different patterns of cognitive activation which will vary in the extent to which they involve physiological, cognitive, behavioural, and other categories of output.

Allowing for (and indeed anticipating) a collection of superordinate devices which between then initiate a broad spectrum of specific activation configurations, arguably aligns the computational approach more easily with the folk conceptions revealed in every-day language use concerning emotions. Acknowledging wide variations in intensity, complexity, and duration makes sense of the proliferation of subtly differing synonyms for emotion states - such as mood, feeling, sentiment, passion, affect, appetite, and so on.

From this computational perspective, if moral elevation is an evolved emotional adaptation, it must manifest as a consistent pattern of cognitive activation which serves to pilot our behaviour through an ancestrally recurrent and complex adaptive challenge. Two ideas have been proposed about the challenge which moral elevation may have been selected to help us overcome. The next section details these proposals, before turning to the predictions they would make about the particular configuration of cognitive activation moral elevation should be expected to generate.

Moral Elevation: Theories About Function

To date, two ideas have been mooted concerning moral elevation's possible adaptive function. Neither have been expounded at any proper length, however, and there have been no empirical tests of predictions derived from either hypothesis. Although their exponents have outlined them only in passing, this section will consider the merits of each proposal; they will be referred to here as the relationship-building hypothesis and the reputation-management hypothesis.

Both ideas start from the same premise: people who behave benevolently make attractive social partners. Forming social relationships (friendships, mate-

ships, teams, alliances, and so on) is integral to our survival and reproductive success, and individuals who have a tendency to cede welfare to others generate more benefits as relationship partners than individuals who do not. Consequently, when we see others behaving benevolently, we see them demonstrating their high value as potential social relationship partners.

The relationship-building hypothesis

Under the relationship-building hypothesis, offered by Haidt and Fredrickson (Haidt, 2000; Fredrickson, 2004), the presence of a high value potential social partner signifies an opportunity. Coming across a “good candidate for cooperation and affiliation” (Haidt, 2000, p.4) presents us with a chance to form a highly fruitful social relationship. The relationship-building hypothesis points out that ancestors who failed to recognise and capitalize on opportunities to build such a crucial resource would have paid a heavy fitness penalty. Haidt and Fredrickson suggest that moral elevation was selected because it helps us to capitalize on these opportunities – they see it as an adaptation which helps us avoid the opportunity cost of failing to form high-quality social relationships (Haidt, 2000; Fredrickson & Branigan, 2001).

According to the relationship-building hypothesis, then, moral elevation is a mental mechanism which kicks in when we see someone else behaving benevolently, allows us to recognise the actor as a good-quality potential relationship partner, and prompts us to take steps which should help us to take advantage of the opportunity this presents. The emotion’s functional output is to motivate and coordinate behaviours which facilitate forming social bonds; “triggering love, admiration, and affection for the altruist, and making affiliative behaviour more likely” (Haidt, 2000, p.4). The other benefitting behaviours which elevation tends to generate operate as social overtures, inviting and increasing the chances of forming relationships with others.

Although the selection pressure that Haidt and Fredrickson point to is undoubtedly real - social relationships are indeed a vital resource, and it is essential that we recognise and capitalize on opportunities to build them – there seems to be a disjunct between the function they propose and the behavioural output that moral elevation produces. The most unique and interesting aspect of the behavioural phenomena associated with elevation – the apparent indiscriminateness of the

benevolent motivation it generates – does not seem to be a plausible way of addressing the pressure identified by the relationship-building hypothesis.

Despite the importance of social relationships, it is also essential that we exercise discretion in choosing whom to build them with. We have limited time and resources, and so we are only ever able to maintain a finite number of affiliations (Tooby & Cosmides, 1996). The choices we make about which individuals to form affiliations with have important fitness consequences, because people differ in their abilities to benefit their social partners. Variations across a host of dimensions (like status, strength, resources, and temperament) mean that relationships with some people bring greater fitness benefits than relationships with others (Noë & Hammerstein, 1994). Individuals who failed to discriminate when making social overtures would quickly end up at a significant fitness disadvantage to individuals who carefully targeted their overtures towards potential partners who would be able to return relatively high benefit yields.

While individuals whose behaviour elicits moral elevation are themselves good candidates for social relationships, and elevated individuals do report motivations to seek a relationship with the person whose behaviour elicited the emotion, by far the most prominent and profound aspect of the experience seems to be a change in how we want to interact with people more generally (Algoe & Haidt, 2009; Haidt, 2000). Elevation seems to make people want to behave benevolently towards others in a remarkably broad sense. This lack of discrimination would make a poor strategy for building an appropriate stock of high-quality cooperative partnerships – at least in the kind of social ecologies that humans inhabit, where the importance of being selective about our partners is high (a difficulty with the relationship-building account which has also been highlighted by (Pohling & Diessner, 2016)).

Because the relationship-building hypothesis has not been expounded at any proper length, it could be that Haidt and Fredrickson intend to suggest that elevation functions to increase the chance of building a relationship exclusively with the benevolent person whose behaviour elicited the emotion. The indiscriminateness of the benevolence that elevation motivates might be thought to function as a signal to the original benevolent actor – a display intended to attract her and increase the chance of being able to build a cooperative relationship with her. But even under

this narrower parsing of moral elevation's putative function, the behaviour the emotion motivates still seems inappropriately indiscriminate for the task.

Signaling one's willingness to behave benevolently so as to attract social partners is an evolutionarily viable strategy, and one that we do seem to engage in successfully (behaviour which will be described in more detail in the next sub-section of this chapter) (Barclay & Willer, 2007; Raihani & Smith, 2015; Roberts, 1998; Van Vugt, Roberts, & Hardy, 2009). It is not, however, the behavioural tendency that we should expect to be dominant in the output of a cognitive mechanism designed to increase our chances of forming a relationship with one particular individual.

If moral elevation was an adaptation for helping us to build a relationship with a particular good candidate for affiliation, this increase in general prosocial motivation should take more of a supporting role alongside stronger motivations to interact specifically with the individual we would like to build a relationship with. Building a single good quality social relationship would be more efficiently, economically, reliably, and precisely achieved with a much more targeted kind of benevolence than that generated by experiences of moral elevation.

The reputation-management hypothesis

Fessler and Haley's reputation-management hypothesis (Fessler & Haley, 2003) points to a different kind of selection pressure created by the presence of high value potential social partners. While it may sometimes represent an opportunity, others' socially attractive behaviour can also present a threat to our relative social value.

Because people vary in their ability to benefit their social partners, and individuals who managed to form relatively high-yield relationships stood at a fitness advantage to those who did not, our ancestors faced a strong selective pressure to form the best quality social relationships they could (Nesse, 2007). The competition created by this pressure resulted in the formation of a biological marketplace (Noë & Hammerstein, 1994), in which each individual bids for access to the most beneficial relationships they can afford. We bid for access by advertising our own value as cooperative interactants; demonstrating our positive partner qualities so as to increase the chances that others will choose to form relationships with us. Through these demonstrations, we create and maintain social reputations, and on the

strength of these reputations we win opportunities to form alliances and bonds with others (Barclay, 2013; Noë & Hammerstein, 1995).

Benevolent behaviour is an excellent way to showcase partner quality - it demonstrates that an individual has both the resources and disposition to benefit anyone they form a social relationship with. Advertising partner quality through benevolence is known as competitive altruism (Roberts, 1998), a form of costly signaling (Zahavi, 1975), and it is thought to explain otherwise puzzling displays of benevolent behaviour across a variety of species (Noë, Van Schaik, & Van Hooff, 1991). Because of its advertising power, individuals can make fitness returns on benevolent behaviour even when the odds of direct reciprocation are low, when it would generate no inclusive fitness benefit, and when there is no threat of coercion. So long as it narrows the odds of us forming high-quality social relationships, by making us attractive in the eyes of others, welfare ceding behaviour can be fitness enhancing (Barclay & Willer, 2007; Fehrler & Przepiorka, 2013; Hardy & Van Vugt, 2006; Nesse, 2007; Van Vugt et al., 2009).

In any market competition, margins are important. To compete, we need to avoid “[lagging] too far behind relative to everyone else” (Alexander, 1987, p105), but advertising comes at a cost. We would lose the fitness advantage that good quality relationships bring if we were to pay significantly over the odds in advertising for them. The level of investment we need to make in order to adaptively manage our social reputations is a fine and moveable line - we need to display a level of benevolence comparable to that demonstrated by others in our social environment (Barclay, 2013; Fessler & Haley, 2003).

When others raise the bar, by engaging in remarkable public benevolence, we face a pressure to increase the investment we make in similar behaviours, so as to safeguard our relative reputational standing. Fessler and Haley suggest that moral elevation is a cognitive device selected to respond to this pressure (Fessler & Haley, 2003). The reputation-management hypothesis proposes that elevation is an adaptation which dynamically increases the investment we make in reputation-management activities, in response to others’ public welfare-ceding behaviours, so as to protect us from the fitness costs of failing to access or maintain adequate quality social relationships.

Unlike the relationship-building hypothesis, the reputation-management hypothesis faces no difficulty explaining the indiscriminateness of moral elevation’s

benevolent motivational output. When it comes to maintaining a reputation for benevolence, what matters is not who we are benevolent to, but that other people see us behaving benevolently. The more indiscriminate we are about our beneficiaries, the greater the chance that our benevolence will be noticed by others and our social reputations will be strengthened¹. If moral elevation is a cognitive device selected to help us manage our social reputations, then the indiscriminate benevolence it elicits would constitute a functional, broadcast signal widening, component of the emotion's output.

Generating predictions about form

Although they point to different selection pressures, these two hypotheses make largely identical predictions about the functional qualities moral elevation should be expected to demonstrate. There are no differences in the predictions either hypothesis would make about the stimuli which should elicit moral elevation (the inputs of the cognitive mechanism). Another individual's observable demonstration of attractively benevolent qualities constitutes both a potential opportunity for a good-quality social relationship and a potential threat to our relative reputational standing.

There are, however, some differences in the predictions each hypothesis would make about the outputs we should expect moral elevation to generate. The most obvious difference has already been discussed – the reputation-management hypothesis would predict that moral elevation should generate a relatively indiscriminate benevolent motivation because this would increase signal broadcast strength, maximizing the efficacy of the behaviour as a reputation-safeguarding device. In addition to this, though, the two hypotheses make different predictions about the motivational priorities moral elevation should produce; a relationship-building adaptation should incline us towards prioritizing opportunities to engage in relationship-building activities, whereas a reputation-management device should incline us towards prioritizing reputation-managing opportunities (a distinction which will be unpacked in Chapter 5).

¹ Although this may only hold within limits – it would likely do little to enhance our social attractiveness if we behaved benevolently towards individuals who presented threats to others, e.g., exploitative or aggressive individuals

Although the majority of the predictions these hypotheses make are not helpful in terms of differentiating their validity, they can help us to explore whether moral elevation shows a consistent pattern of activation indicative of special design. The following chapters investigate whether moral elevation displays traits which conform to the expectations of either theory of function, as seen through the lens of a computational model of emotions. Figure 1 illustrates the basic computational form moral elevation should be expected to display if it functions as an adaptive superordinate mechanism-orchestration device (with labels indicating which components of the model will be explored in which empirical sections of the thesis).

The investigations which follow are offered as tentative, exploratory steps towards investigating the possibility that moral elevation may operate in the manner described by this model. The thesis is organized into two parts, intended to reflect the computational structure we should expect from this kind of cognitive adaptation. The first part looks at what characteristics the 'input' side of moral elevation should be expected to demonstrate if it is a functional mental mechanism, and the second part considers what characteristics the 'output' side should be expected to demonstrate.

Chapter 2 offers a review of the literature concerning moral elevation's elicitors, and outline predictions about elevation's inputs which are grounded in a functionalist analysis. Specifically, the chapter considers what specific features of others' social behaviours should be expected to elicit moral elevation, and whether any other exogenous or endogenous variables might moderate the relationship between others' social behaviours and the elicitation of moral elevation. Chapters 3 and 4 report the results of studies intended to test these predictions.

Chapter 5 offers a review of moral elevation's behavioural and motivational effects, and present functionalist predictions about elevation's observable outputs. The chapter focuses in particular on the possible behavioural and motivational outputs of moral elevation, though other sub-components which have shown to or may be affected by elevation are illustrated in Figure 1 (marked with dashed lines).

Chapter 6 reports the results of studies intended to test these predictions.

Chapter 7 discusses the results of all studies reported here, draws conclusions, and offers suggestions for future research directions.

Figure 1: Computational Model of Moral Elevation

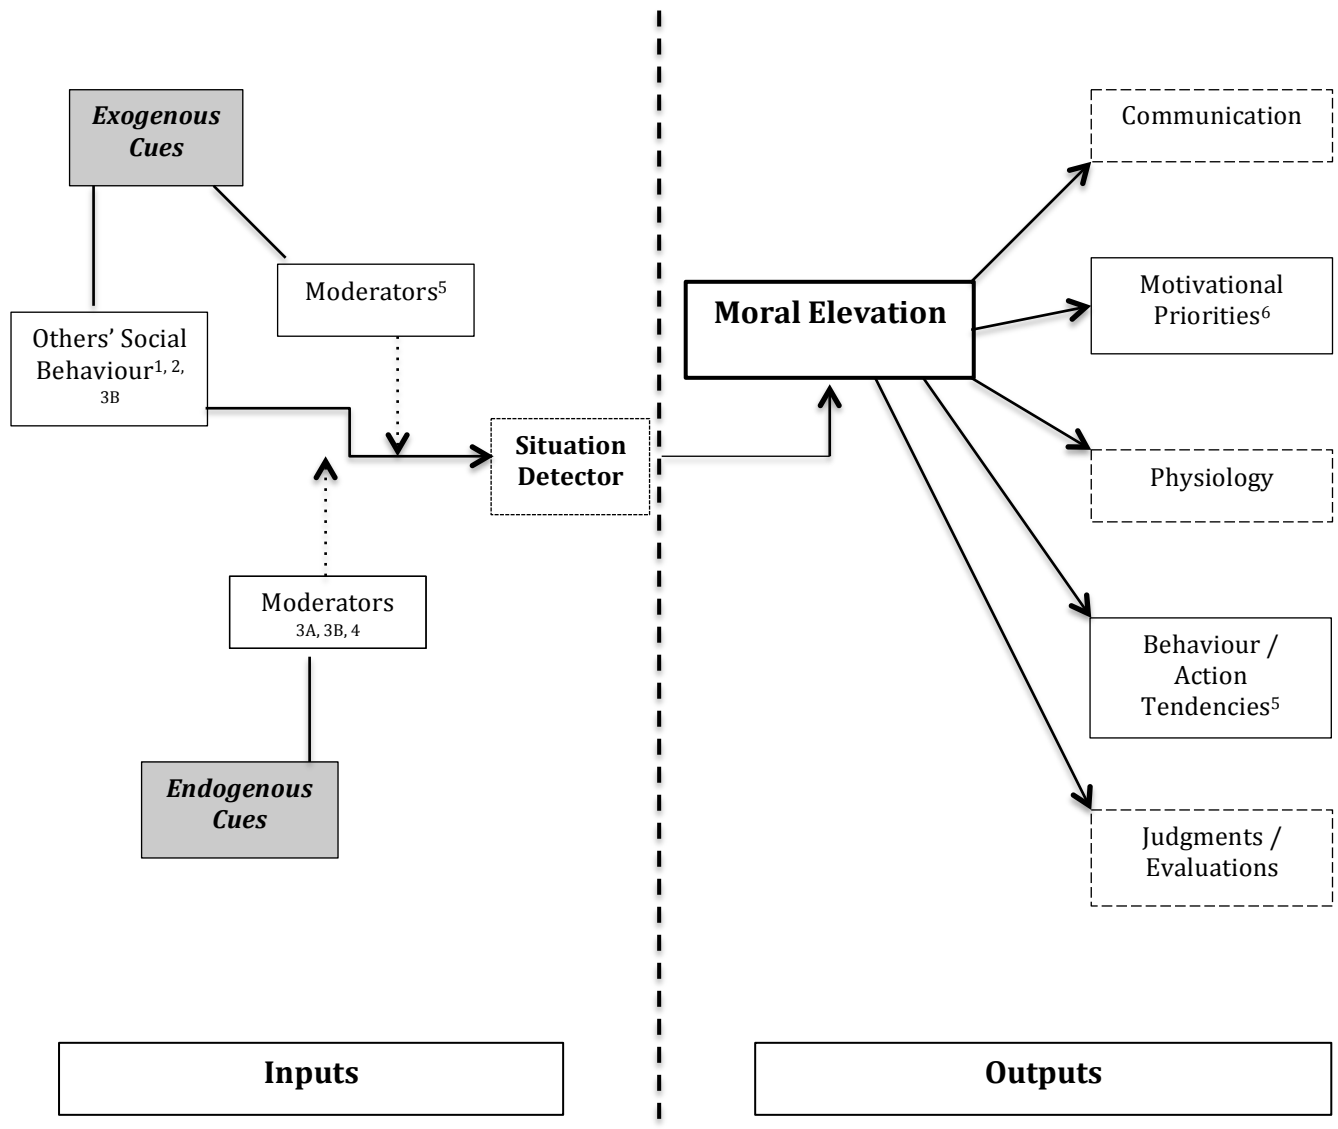


Figure 1. Model showing the proposed basic computational structure of moral elevation, arranged as inputs and outputs.

Note: ¹ = Component explored in Study 1; ² = Component explored in Study 2; ^{3A} = Component explored in Study 3A; ^{3B} = Component explored in Study 3B; ⁴ = Component explored in Study 4; ⁵ = Component explored in Study 5; ⁶ = Component explored in Study 6.

Chapter 2: Moral Elevation's Elicitors

Abstract

This chapter begins with a summary of empirical findings offered in the literature to date which relate to the eliciting conditions for moral elevation. Four predictions about elevation's elicitors are then derived from the evolutionary hypothesis outlined in the previous chapter. The first two predictions concern characteristics of an elevating actor's behaviour; specifically, whether the behaviour is benefit conferring or simply moral norm-adhering, and whether or not an elevating act is costly for the actor to perform. The second two predictions concern characteristics of elevating actors themselves; specifically, whether or not an actor has any obvious reasons to value the recipient of their benevolent actions, and whether or not an actor has any other socially attractive partner qualities.

Literature Review

The moral elevation literature has so far been preoccupied overwhelmingly with the emotion's motivational output, and there has been comparatively little attention paid to describing the types of event which elicit it. As the chosen name for the emotion suggests, moral elevation is most often described as being elicited by morally praiseworthy behaviour. Although the full scope of moral behaviour encompasses more than just benevolence (as Haidt himself points out elsewhere; Haidt, Koller, & Dias, 1993), the qualifiers and examples used throughout the literature make clear the assumption that moral elevation is a reaction to seeing people increasing the welfare of others. 'Morally good' is often used interchangeably with terms like 'kind' (Freeman et al., 2009; Haidt, 2000, 2003b; Landis et al., 2009; Schnall & Roper, 2012; Silvers & Haidt, 2008), altruistic (Ellithorpe, Ewoldsen, & Oliver, 2014; Haidt, 2000, 2003a; Piper, Saslow, & Saturn, 2015), and helpful (Haidt, 2003a; Thomson et al., 2014; Thomson & Siegel, 2013).

In early research, when participants were asked to recall incidents which had elicited the sensations associated with moral elevation they "most commonly cited circumstances that ... involved seeing someone else giving help or aid to a person who was poor or sick, or stranded in a difficult situation" (Haidt, 2000., p.2). Vignettes used by researchers to provoke moral elevation in experimental work since have described behaviours such as mentoring (Schnall & Roper, 2012; Schnall et al., 2010; Silvers & Haidt, 2008), helping the homeless (Algoe & Haidt, 2009; Aquino, McFerran, & Laven, 2011), making charitable donations (Aquino et al., 2011), and life-saving heroism (Englander, Haidt, & Morris, 2012).

There is agreement in the literature that not all acts of benevolence are likely to elicit moral elevation, but there have been few attempts to specify closely what features of any particular act would predict its capacity to arouse the emotion. The qualifications offered by many authors are exceptionally vague, for example that an action must be 'morally excellent', 'morally exemplary', or 'morally beautiful' in order to elicit moral elevation (Haidt, 2000, 2003a; Haidt, 2003b; Monin, 2007; Algoe & Haidt, 2009; Landis et al., 2009; Vianello et al., 2010; Cox, 2010; Schnall et al., 2010; Strohminger, Lewis, & Meyer, 2011; Aquino et al., 2011; Englander et al., 2012; Schnall & Roper, 2012; Lai et al., 2014; Siegel, Thomson, & Navarro, 2014; Thomson et al., 2014). Other papers suggest benevolence which is unexpected, extraordinary, uncommon, or profound (Freeman et al., 2009; Haidt, 2000; Piper et

al., 2015). A summary of terms used to describe elevation's elicitors is provided in Table 1 below.

Only one paper has focused its investigations on features which might moderate the relationship between benevolent behaviour and moral elevation. Thompson and Siegel (2013) found that participants who read a vignette about someone helping a person of good character reported stronger feelings of moral elevation, and subsequently made higher donations to a charity (taken to be a behavioural measure of elevation), than did participants who read about someone helping a person of bad character. They also found that participants who were asked to imagine someone doing a good deed at great expense made higher donations to charity (although they didn't report stronger feelings of elevation) than did participants who were asked to imagine someone doing a good deed at little expense. Beyond these findings, there is as yet no information available about which benevolent scenarios are most likely to elicit moral elevation or which would elicit it more strongly than others.

None of the descriptions of moral elevation's elicitors offered in the literature, and neither of Thompson and Siegel's predictions, are grounded in the logic of a functionalist analysis. The sections that follow will discuss what characteristics we should expect to see in the situation detecting component (the 'input side') of moral elevation if the emotion has been shaped by natural selection to pick up on others' displays of high social partner value. The purpose of this chapter is to make predictions about what cues moral elevation should be sensitive to by considering what reliable, observable features of a situation tell us that someone else is likely to be a high-quality social affiliate. Testing these predictions will help us to evaluate the likelihood that moral elevation is an adaptation; positive results would indicate a higher probability that the emotion is a functional feature of our species-typical cognitive architecture.

Table 1: Descriptions Used for Behaviour Which Elicits Moral Elevation

Descriptor	Publications
Moral	Landis et al., 2009; Cox, 2010; Thomson & Siegel, 2013; Thomson et al., 2014; Siegel et al., 2014
Morally beautiful	Haidt, 2000; Haidt, 2003a; Haidt, 2003b; Algoe & Haidt, 2009; Landis et al., 2009; Vianello et al., 2010; Cox, 2010; Strohminger et al., 2011; Englander et al., 2012; Thomson & Siegel, 2013; Lai et al., 2014; Siegel et al., 2014; Piper et al., 2015
Morally admirable	Haidt, 2000
Morally exemplary	Monin, 2007; Algoe & Haidt, 2009; Freeman et al., 2009; Cox, 2010; Aquino et al., 2011
Morally uplifting	Algoe & Haidt, 2009
Morally excellent	Freeman et al., 2009; Landis et al., 2009; Vianello et al., 2010; Schnall et al., 2010; Schnall & Roper, 2012; Thomson et al., 2014
Morally good	Freeman et al., 2009; Landis et al., 2009; Cox, 2010; Aquino et al., 2011
Morally virtuous	Freeman et al., 2009; Aquino et al., 2011; Englander et al., 2012; Ellithorpe et al., 2014; Thomson et al., 2014; Siegel et al., 2014; Van de Vyver & Abrams, 2015
Virtuous	Haidt, 2000; Silvers & Haidt, 2008; Algoe & Haidt, 2009; Freeman et al., 2009; Landis et al., 2009; Vianello et al., 2010; Cox, 2010; Thomson & Siegel, 2013; Erickson & Abelson, 2012; Thomson et al., 2014; Piper et al., 2015; Van de Vyver & Abrams, 2015
Exhibitions of humanity's better nature	Haidt, 2000; Aquino et al., 2011; Freeman et al., 2009; Monin, 2007
Saintly	Haidt, 2000; Haidt, 2003b
Good	Haidt, 2000; Silvers & Haidt, 2008; Algoe & Haidt, 2009; Schnall et al., 2010
Superhuman	Haidt, 2003a
Exhibitions of strength of character	Landis et al., 2009
Blurring human/God divide	Haidt, 2003a

Four predictions will be put forward in this chapter. Predictions 1 and 2 concern characteristics of morally elevating acts themselves, and predictions 3 and 4 concern characteristics of the people involved in morally elevating acts.

Morality vs. Benevolence

As outlined in the literature review above, there seems to be an implicit assumption amongst authors that 'moral elevation' may be something of a misnomer. Although there is a vast overlap between benevolence and morality (more often than not, behaviours which are benevolent are also considered morally praiseworthy), the two categories are fully dissociable (Kurzban, DeScioli, & Fein, 2012). Some behaviours which involve increasing the welfare of another (e.g. helping a wanted murderer to evade capture) would almost certainly be considered immoral, and some behaviours which are generally considered moral (e.g. observing religious food taboos) are, at the least, not obviously benevolent. Thompson and Siegel (2013) describe results which suggest that benevolent but morally questionable behaviour is less likely to elicit moral elevation than benevolent and morally laudable behaviour, but it remains an intuitive assumption that moral behaviour which does not involve conferring benefits on others is also less likely to elicit moral elevation (see, for example, Landis et al., 2009). Because it has not been explicitly explored, it remains an open question whether moral behaviour is sufficient to elicit moral elevation, or whether benevolence is also necessary.

From the functionalist perspective of either the relationship-building hypothesis or the reputation-management hypothesis, it makes sense to expect that moral elevation would be more sensitive to benevolent behaviour than to moral behaviour per se, because benevolent behaviour is a better indicator of partner quality than simply moral behaviour. Non-benevolent moral behaviours may be indicative of traits which would bring distal benefits to an individual's relationship partners, e.g., individuals who adhere to purity norms might enjoy the health benefits of good hygiene, decreasing the odds that any social relationship with them would be cut short by illness. But advantages of this ilk pale in comparison to the direct benefits that come from an individual who undertakes to deliberately increase others' welfare. If moral elevation is an adaptation selected to react when others display attractive partner qualities, then we should expect it to be much more readily

responsive to others' benefit-conferring behaviours than to their norm-adhering behaviours.

Prediction 1: Moral elevation is more likely to be elicited by others' benevolent behaviour than by behaviours which are morally praiseworthy but non-benevolent

This prediction is yet to be explored. There is a uniform tendency in the elevation literature to simply equate morality with benevolence - a conflation well evidenced in the prompts which are often used when asking participants to recall elevating incidents from their past, e.g., "think of a time when you saw someone do something wonderful, a very good deed, **to someone** else, but not to you" (Haidt, 2003a, emphasis added. See also Algoe & Haidt, 2009; Aquino et al., 2011; Cox, 2010; Siegel et al., 2014; Thomson & Siegel, 2013). As a consequence, there has been no deliberate demonstration of the prediction being met.

The results of one study by Algoe and Haidt (2009) offer a little support for the idea that elevation is more likely to be elicited by benevolent behaviour than by non-benevolent but otherwise moral behaviour. Participants who were asked to recall witnessing a benevolent act reported having affective and motivational responses to the act which fit the description of moral elevation, whereas participants who were asked to recall witnessing someone overcoming adversity (a feat which demonstrates a number of non-benevolent moral virtues) reported having affective and motivational responses which fit the description of admiration/inspiration. Because Algoe and Haidt had implicitly defined moral behaviour as benevolent and non-benevolent behaviour as non-moral, they classified elevation as a reaction to moral excellence.

Benefit Delivery: Ability vs. Willingness

The size of the benefit one person delivers to another is determined by two factors – the actor's ability to deliver benefits (the scale of the resources she has available), and her willingness to do so (her generosity). For Jenny to benefit Juan, she must both have resources to offer, and be motivated to offer them. Recent research has demonstrated that, when it comes to gauging others' value as social partners, we pay much greater attention to their generosity than to their ability to confer benefits. Access to resources varied widely ancestrally, as a result of

unpredictable factors such as poor luck, injury, and illness. An individual's willingness to deliver benefits to others, on the other hand, tends to remain relatively stable over time. Consequently, an individual's generosity, and not the resources she had available, made the best predictor of the likely long-term benefits she would offer as a social partner (Delton & Robertson, 2012). Pressure to pick the best available social partners led to selection for cognitive mechanisms which allow us to accurately gauge others' generosity, and use this gauged value in estimating how profitable specific relationships would likely be, relative to others.

Generosity is a factor of how much an individual values another person's welfare relative to her own. The more she values the welfare of the other, the more motivated she will be to trade-off her own welfare in order to increase theirs. We can describe this relative valuation as a ratio – one individual's welfare trade-off ratio (WTR) towards another (Delton, 2010). If Jenny values Juan's welfare half as much as her own, then her WTR_{Juan} would be 1:2, or 0.5; if she values Juan's welfare just as much as her own, her WTR_{Juan} would be 1; if she did not value Juan's welfare at all, then WTR_{Juan} would be 0; and so on. Our minds compute and store a WTR for each specific other in our social environment. When we face a choice between behavioural alternatives with opposing fitness consequences for ourselves and an other, the WTR we hold for that other stands as one of three parameters in the decision rule which guides our behaviour. The other two parameters are the costs (C) and benefits (B) which would result from a particular choice, and the decision rule runs as follows (Delton, 2010; Tooby & Cosmides, 2008; Tooby et al., 2008):

- i) For actions which would benefit the other at a cost to the self, take the action if and only if $C_{Self} < B_{Other} \times WTR_{Other}$
- ii) For actions which would benefit the self at a cost to the other, take the action if and only if $B_{Self} > C_{Other} \times WTR_{Other}$

When we have an opportunity to benefit another individual, we take the opportunity so long as the cost we would incur is lower than the benefit the individual would derive, discounted by our valuation of their welfare. This decision rule sits at the root of our ability to adaptively regulate our welfare-ceding behaviour.

In order to differentiate, then, between an individual's ability to benefit another and her willingness to do so, we need to pay attention to the cost : benefit structure of her behaviours. Looking at benefit size in isolation does not allow us to distinguish between the well-resourced-but-mean and the poorly-resourced-but-generous. For example, if we only knew that Jess had increased Juan's welfare by 50 'fitness points', and Jenny had increased Juan's welfare by 10 'points', Jess may seem like the more generous of the two. If we know the proportionate costs they each incurred in increasing Juan's welfare, however, we can infer (a lower bound for) the WTR they each hold towards Juan. If Jess's gift to Juan cost her 10 'points', while Jenny's cost her 5, then we could deduce that Jess has a minimum WTR_{Juan} of 0.2 (10/50) whereas Jenny has a minimum WTR_{Juan} of 0.5 (5/10). Although Jess delivered a larger benefit to Juan (and incurred a greater absolute cost), Jenny incurred a higher proportionate cost than Jess, demonstrating a higher valuation of Juan's welfare – Jenny is more generous than Jess, and is likely to have a greater positive impact on Juan's welfare over time.

Because of this predictive utility, we use information about the cost : benefit structure of others' actions when we make judgments about which social partners to pick and invest in relationship with. Rather than simply preferring people who deliver the largest benefits to others, we pay close attention to the proportionate costs incurred in others' benevolent behaviours. Social partners who incur higher proportionate costs to deliver benefits to others are viewed as preferable to partners who incur lower proportionate costs, even if the size of the benefits they deliver is larger (Bliege Bird & Power, 2015; Delton et al., 2012; Delton & Robertson, 2012; Lim, 2012; Sell, 2005).

If moral elevation is a cognitive mechanism which has been naturally selected to respond to indications of high partner value in others, then we should expect it to be sensitive to information which we are already known to track and use in estimating partner value. The situation detecting component of moral elevation should be functionally integrated with our 'welfare trade-off ratio estimator'. If this were the case, then moral elevation would be more readily or more strongly activated by benevolent behaviours which require the actor to incur relatively high proportionate costs.

Prediction 2: A benevolent act will elicit stronger feelings of moral elevation if the actor has to pay a higher proportionate cost to perform the act

Thompson and Siegel (2013) found that participants who were asked to imagine benevolent acts which were very costly to perform reported stronger feelings of elevation than did participants who were asked to imagine less costly benevolent acts. However, because the authors did not control for the size of the benefits produced by these benevolent acts, it is not possible to draw conclusions from their results about whether the proportionate cost incurred by the actor has an effect on the potential of an act to elicit moral elevation.

Incentives to Trade-off Welfare

There are a number of incentives for individuals to trade-off their welfare for specific others, including inclusive fitness incentives, the promise of direct or indirect reciprocity, and avoiding costly conflicts over resources. Variables like kinship, likelihood of reciprocity, and relative formidability all affect how willing we are to cede welfare to others (Barclay, 2004; Curry et al., 2013; Hardy & Van Vugt, 2006; Jensen & Petersen, 2011; Sell et al., 2008; Smith, Kish, & Crawford, 1987). Our minds use a myriad of cues to gauge the magnitude of each variable (Barclay, 2013; Delton & Robertson, 2012; Lieberman et al., 2007; Sell et al., 2009), and these magnitudes are combined in forming our welfare trade-off ratio toward any given other (Delton, 2010; Lim, 2012).

The presence and strength of these incentives to benefit others will be randomly distributed across the social interactions that we get to observe (only a fraction of those that go on between those around us). As a result, they can produce confounds when we try to infer how much an individual would benefit us as a social partner by observing how much she benefits others. Just because Jenny benefits Juan greatly, we can not safely assume that she would benefit us equally well if Juan happens to be her son, her husband, or her boss. In order to accurately appraise others' value as potential social partners, our minds control for the presence of strong incentives in the specific episodes of benevolence we get to observe. We judge others as comparatively less socially attractive ('altruistic', 'kind', 'nice', and 'generous') when they benefit from their own benevolent acts (Lin-Healy & Small, 2013). It should be assumed, then, that benefit conferring behaviour will have less

power to elicit moral elevation when there are clear incentives for the actor to engage in it.

Prediction 3: A benevolent act will elicit stronger feelings of moral elevation if there are no obvious incentives for the actor to value the welfare of the recipient

Indirect Benefit Delivery

Finally, a tendency to behave benevolently is not the only indicator that an individual may be a high-value social relationship partner. Variables like strength, skill, health, status, intelligence, and attractiveness can affect how able another individual would be to access and retain resources, and thereby would affect their likely long-term ability to deliver benefits. Individuals who are well-endowed on any of these other dimensions may also generate positive externalities (indirect benefits produced by proximity with the individual) which would benefit their social relationship partners (Barclay, 2013). For example, associating with males who are strong enough to deter ambushers would reduce one's own chance of being ambushed.

Although, as argued above, non-benevolent attractive partner qualities are expected to be insufficient to elicit moral elevation (benefit conferral should be a necessary condition), other attractive traits may still be expected to have a compounding effect on the ability of benevolent acts to elicit elevation. A benevolent actor who already commands a high partner-value, by virtue of one of these other traits, would present a more attractive potential relationship prospect to onlookers, and therefore also a greater threat to the relative market value of others, than a benevolent actor of otherwise relatively low partner-value.

Prediction 4: A benevolent act will elicit stronger feelings of moral elevation when the actor has other traits indicative of high partner-value

For ease of presentation, tests of these four predictions will be presented across two chapters. Chapter 3 will report the results of studies intended to test predictions 1 and 2; predictions about which characteristics of others' behaviour are likely to elicit moral elevation. Chapter 4 will report the results of studies intended to test predictions 3 and 4; predictions about recipient and actor characteristics which

may moderate the ability of others' behaviour to elicit moral elevation (the studies presented in Chapter 4 will also offer additional tests of predictions 1 and 2). There will be a general discussion of empirical results relating to all four predictions about moral elevation's elicitors at the close of Chapter 4.

Chapter 3: Empirical Tests of Predictions 1 and 2

Abstract

This chapter reports the results of two studies. Study 1 aimed to test the prediction that benefit conferral is a necessary condition for another's behaviour to elicit moral elevation, and Study 2 aimed to test the prediction that the power of a benevolent act to elicit moral elevation will increase in line with the proportionate cost incurred by the benevolent actor. The results of Study 1 failed to provide support for the prediction – implications and potential study design limitations are discussed. The results of Study 2, in contrast, provided good support for the hypothesis.

Study 1: Is Benefit Conferral Necessary to Elicit Moral Elevation?

Study 1 aimed to test prediction 1:

Prediction 1: Moral elevation is more likely to be elicited by others' benevolent behaviour than by behaviours which are morally praiseworthy but non-benevolent

Design

Study 1 employed a between-subjects experimental design. Participants were randomly divided into two conditions: participants in condition 1 were asked to read a vignette describing a benefit conferring act of self-sacrifice, and those in condition 2 were asked to read a vignette describing a non-benefit conferring but morally laudable act of self-sacrifice. Participants then completed a moral elevation measure; it was anticipated, in line with prediction 1, that participants in condition 1 would report stronger feelings of moral elevation than participants in condition 2.

In order to explore the cross-cultural consistency of moral elevation, study 1 recruited participants from two geographical regions: India, and the USA. With the exception of one small-sample and informally reported study (Haidt, 2003a), no prior research on moral elevation has included participants from non-Western cultures, and so it remains an open question whether experiences of moral elevation vary across social and ecological circumstances.

It should be noted that participants in study 1 were also asked, in the final stage of the study, to complete a 9-item questionnaire measure intended to test prediction 6 of this thesis. This was done in order to maximize data-collection efficiency, but for ease of exposition the results from that questionnaire measure will be reported and discussed in Chapter 6, after prediction 6 has been put forward (the sample characteristics will be reiterated at that point).

Method

Participants

208 participants (112 from the USA and 96 from India) were recruited via Amazon Mechanical Turk between 24-25/02/2019, and were paid \$1 for their time.

Ethical approval for the study was granted by Brunel University London's College of Health and Life Sciences Research Ethics Committee on 12/12/2018.

Participants' age range was 18-68 ($M = 34.70$, $SD = 10.02$), and 44.8% were female. 64.2% of participants reported having some kind of religious affiliation, with the remaining 35.8% explicitly identifying as non-religious.

Participants were screened using the 3 attention checks planted at different points in the survey. Directly after reading the short story, participants were asked two basic comprehension questions about the text itself, and the very end of the survey, participants were presented with a free-text comment box, and asked to answer honestly whether they had been paying attention during the study. Participants who got either of the first two questions wrong, or who answered anything other than affirmatively to the last question, were excluded from all analyses. The remaining N was 134 (37 from India and 97 from the USA).

Measures

Both vignettes used (the morally elevating version and the control) are given in full in Appendix A. Both vignettes describe a light-aircraft, carrying two men, Oshin and Mika, which crashes in a remote location. In the *Benevolent* version of the story (morally elevating condition), both men survive, but Mika is badly injured. There are limited rations, and Oshin opts to starve in order to feed the rations to Mika. In the *Moral* version (control condition), only Oshin survives the crash – Mika is killed outright. There are no rations whatsoever, but Oshin opts to starve rather than resort to cannibalism. The text of each story was kept as near to identical as possible. Both stories describe Oshin's struggle with hunger, and the sacrifice he makes in order to keep Mika alive / avoid cannibalism. Cannibalism was chosen for the non-benevolent moral condition because it is a widely-held moral taboo which does not centre on delivering a benefit to another individual.

Moral elevation was measured using a modified version of a scale developed by Vianello and colleagues (Vianello et al., 2010). Vianello et al.'s 8-item questionnaire taps the affective, physical, and motivational aspects of the experience of moral elevation. In full the items run: 'I feel more open towards others'; 'I feel like I'm a better person'; 'I feel warmth in my chest'; 'I feel a lump in my throat'; 'my muscles are relaxed'; 'I would like to do something good for other people'; and 'I would like to behave as [story protagonist] behaved'. The scale was developed by

Vianello and colleagues using the results of an exploratory factor analysis performed by Algoe and Haidt (2009), and its three factor structure was supported by Vianello et al.'s own confirmatory factor analysis. Vianello et al. reported high scale reliability; $\alpha = 0.82$. This measure was adapted here to an ipsative, or forced choice, format. Each of the 8 items was paired with two other answer choices; a positively-valenced sensation not associated with moral elevation, and a 'neither' option. For example:

Item 1) *As I read the story I felt:*

- a) More open towards others (*Elevation option*)
- b) Like a strong and capable person
- c) Neither of the above

The full ipsative measure is included in Appendix A.

The order of presentation for the elevation and control response options in each item was counterbalanced across the 8 questions. The task instructions given with the measure ran as follows:

"Please read each of the following pairs of statements. For each pair, please pick the option which best describes how you felt whilst you read the story. If you felt both of the options listed, please pick the option which you felt most strongly. If you felt neither of the sensations listed, please pick the third option, 'neither'."

A forced-choice format was used for the moral elevation probe to introduce methodological variety (the original 8-item Likert-type scale is used in studies 4 and 5 of this thesis). The advantage of having two 'positive' sounding response options for each item is that it minimizes any possible social-desirability bias in participants' answers. Participants who may otherwise over-report their experiences of moral elevation in order to avoid using the negative or lower end of a numeric scale are given the opportunity to choose a different positive response or a null response instead.

Participants were also asked to provide basic demographic information; age, sex, and religious affiliation.

Procedure

Participants were directed via hyperlink to a survey hosted on Qualtrics, described as a psychological study looking at emotional reactions to a short story, and were split randomly by Qualtrics into either condition 1 (experimental) or condition 2 (control).

Participants' moral elevation scores were calculated by giving a score of 1 to each response option from the 8-item elevation scale, and a score of 0 to each control response option and to each 'neither' response option. This gave a minimum score of 0 and a maximum score of 8, with higher scores representing stronger feelings of moral elevation. Data analysis was conducted using IBM's SPSS statistics package.

Results

Table 2 shows the mean moral elevation scores for male and female participants, religious and non-religious participants, and participants from each country and in each condition.

Table 2

Moral Elevation by Country, Condition, Sex, and Religious Affiliation

Group	<i>n</i>	<i>M(SD)</i>
Country		
India	37	4.05 (1.53)
USA	97	3.13 (2.05)
Condition		
Benevolent	65	3.25 (2.03)
Moral	69	3.26 (1.89)
Sex		
Male	74	2.93 (1.97)
Female	60	3.95 (1.81)
Religious		
Yes	86	3.74 (1.78)
No	48	2.75 (2.12)

Incomplete data (i.e. an insufficient number of participants at each possible moral elevation score) prevented the use of an ordinal regression model. Although moral elevation was significantly non-normally distributed ($D(208) = .15, p < .001$), a linear regression model was explored. A backwards step-wise multiple linear regression analysis was conducted, using condition, country, sex, and religiosity as predictors, and moral elevation as the dependent variable. The resulting coefficients are summarized in Table 3.

Table 3
Regression Coefficients: Condition, Country, Sex, and Religiosity as Predictors of Moral Elevation

Predictors	B	SE B	β
Step 1			
Constant	3.30	.60	
Country: India vs. USA	-.24	.44	-.05
Religious: No vs. Yes	.77	.39	.19*
Sex: Male vs. Female	-.72	.42	-.19*
Condition: Control vs. Elevated	.27	.34	.07
Step 2			
Constant	3.05	.39	
Religious: No vs. Yes	.86	.35	.22*
Sex: Male vs. Female	-.72	.33	-.18*
Condition: Control vs. Elevated	.28	.35	.07
Step 3			
Constant	3.20	.35	
Religious: No vs. Yes	.86	.35	.22*
Sex: Male vs. Female	-.72	.34	-.18*

Note: R^2 for Step 1 = .10 ($p = .02$), ΔR^2 for Step 2 = <.01 ($p = .60$), ΔR^2 for Step 3 = 0.1 ($p = .40$). * $p < .05$.

Sex and religiosity were the only variable to emerge as significant predictors of moral elevation, accounting for 9% of the variance in moral elevation combined ($F(2,122) = 5.97, p = .003$). As illustrated in Table 2, the mean elevation scores of female participants were higher than those of male participants - a difference which was statistically significant; $U = 3981.5, p = .003$. The moral elevation scores of

religious participants were also significantly higher than those of non-religious participants; $U = 5138.5, p = .02$.

Neither condition nor country predicted moral elevation. The failure of condition to affect moral elevation was confirmed by a Mann-Whitney U test; $U = 5715, p = .47$, upper CI = $-.95$, lower CI = $.42$. Mean moral elevation scores for participants in each condition are shown in Figure 2.

Figure 2. Moral Elevation by Condition




Figure 2. Mean moral elevation scores for participants in each story condition. Error bars represent 95% confidence intervals.

Discussion

Study 1 found no support for prediction 1; the story condition into which participants were placed failed to predict moral elevation scores, and there was no significant difference in reported moral elevation between participants who read the benevolent story and participants who read the morally praiseworthy but non-benevolent story. The relatively wide confidence interval of the t-test suggest that the study may have lacked the power to find the predicted effect. In addition, no manipulation check was included to make sure that participants reading the control story did not view the protagonist's behaviour as benefit conferring, and so it may have been that, even though Mika was dead, participants viewed Oshin's behaviour as beneficial to Mika. The fact that participants with a religious affiliation reported significantly stronger feelings of moral elevation than non-religious participants leaves open the possibility that some idea of an after-life or belief in the supernatural may have prompted some participants to view Mika as a beneficiary, despite him being un-benefit-able in a practical sense. This methodological error (the failure to include a manipulation check) precludes drawing any firm conclusions from study 1 about the sensitivity of moral elevation to benevolent vs. non-benevolent but morally praiseworthy behaviours.

Although there was no support for prediction 1, some of the secondary variables measured in study 1 did generate results of interest. There was some evidence that moral elevation is not an emotional phenomenon isolated to Western cultures; participants from India (a hitherto un-represented cultural group in elevation research) reported just as many sensations characteristic of moral elevation than did participants from the USA (the staple cultural sampling pool). Additionally, female participants and religiously inclined participants reported significantly stronger feelings of moral elevation than male participants and non-religious participants. These results fit with trends reported (though never noted as features of interest) in the wider moral elevation literature (e.g. (Aquino et al., 2011); (Algoe & Haidt, 2009); and (Landis et al., 2009). Because tests of sex-differences are not routinely reported in the extant moral elevation literature, and tests of religiosity-differences even less so, the robustness and magnitude of these trends are as yet difficult to gauge – the remaining studies in this thesis will provide opportunities to gather more information about these relationships.

Study 2: Is Moral Elevation Sensitive To Proportionate Cost Incurred By The Actor?

Study 2 aimed to test prediction 2:

Prediction 2: A benevolent act will elicit stronger feelings of moral elevation if the actor has to pay a higher proportionate cost to perform the act

Design

Study 2 employed a between-subjects experimental design. Participants were asked to rate how morally elevated they felt in response to a short story about a benevolent act which had required the benefactor to incur either a low, moderate, or high cost. In line with prediction 2, it was expected that participants who read about the high cost benevolence would report stronger feelings of moral elevation than participants who read about the low cost benevolence.

Method

Participants

130 participants were recruited between 03-05/05/2013 using the online crowdsourcing platform Crowd Flower, and were paid \$0.25 for their time. Ethical approval was granted for the study on 15/03/2013.

The age range of participants was 18-81 ($M = 33.36$, $SD = 12.16$), and 48.5% were female. 61.5% of participants indicated that they held a religious affiliation. 79.1% of participants were from the USA, 20% from India, and 0.8% each were from the UK, New Zealand, and Uruguay.

Measures

All participants read the same vignette describing an altruistic act – the vignette only described the protagonist's behaviour and the benefit it produced, but the cost that the protagonist incurred was not mentioned. In full, the passage ran:

Alex is running late, so he takes a shortcut through the park. He's walking briskly and is fairly sure he can make it on time when he comes across an elderly lady

wandering across the grass. She's clearly confused and a bit agitated. Alex pauses to ask her what's wrong, and it quickly becomes apparent that the lady is suffering from early stage dementia and doesn't know how to find her way back to her home. Alex spots a policeman nearby and gently takes the lady over to him. He explains the situation, and the policeman assures Alex that he'll take care of the lady and see that she gets home OK. As Alex turns to leave, the elderly lady takes his arm and pleads with him to stay; she's obviously distressed and seems to find Alex's presence comforting. Alex glances at his watch; if he stays, he'll miss it for sure. He thinks about the lady's distress and comes to a decision - 'Of course I'm staying, don't worry' Alex says to her. He and the policeman manage to work out where the lady lives, and return her to her carers unharmed.

Participants rated their experience of moral elevation in response to this passage on a 7 point Likert scale anchored at 'Not at all elevated' (1), 'Somewhat elevated' (4), and 'Very elevated' (7). Because moral elevation is unlikely to be known by that moniker, participants had already been asked to read a short description of moral elevation and its effects – the descriptive passage ran:

This is a study about an emotion known as 'moral elevation'. Moral elevation is what people often feel in response to seeing others act as moral exemplars. An important part of what causes moral elevation is that the good deed you witnessed is done to someone other than yourself – that's what sets it apart from gratitude. Think about how you feel when you see someone do something really good, kind, selfless, or virtuous for another person (not yourself). People commonly use phrases like 'touched', 'moved', 'inspired', or 'morally uplifting' to describe the feeling of moral elevation, and report a warm sensation in the chest and a desire to do good deeds themselves as a result of the emotion. Morally elevating actions show us the better side of human nature, and make us want to try to be better people ourselves. In what follows, we will use the word 'elevated' to refer to this emotion.

After the initial vignette and moral elevation probe, participants read a short passage describing the cost that the protagonist had incurred in order to deliver the benefit. Each participant saw either the high cost, moderate cost, or low cost version of the passage:

High cost condition: Because he stayed with the elderly lady, Alex missed his flight to Paris. He had saved for months to visit the place where his parents met. Alex suffers from a terminal illness and has been given a prognosis of 6 months; even if he could raise the funds, he would not be able to reschedule his trip before he becomes too ill to travel.

Moderate cost condition: Because he stayed with the elderly lady, Alex missed his interview for a promotion at work. He would now have to hope that his managers understood and would reschedule, or else wait for another opportunity for promotion to come up.

Low cost condition: Because he stayed with the elderly lady, Alex missed getting his car to the garage to be waxed that day. He would have to ring the garage and re-arrange the appointment for tomorrow.

Finally, participants were asked to rate, again on a 7 point Likert scale, whether there was any change in the moral elevation they felt after learning about the cost that the protagonist incurred; the scale was anchored at 'Much less' (1), 'The same' (4), and 'Much more' (7).

Procedure

Participants were directed via hyperlink to a questionnaire hosted on Survey Monkey. They were informed that the study was interested in how morally elevating different acts are, and were asked to read the short descriptive passage explaining moral elevation. Participants were randomly assigned by Survey Monkey to one of three conditions (low cost, moderate cost, and high cost). Each participant read the initial vignette, and rated their experiences of moral elevation from 1-7. Participants were then shown one of the passages describing the cost the altruist had incurred, dependent on the condition in which they had been placed, and were asked to rate how much more or less morally elevated they now felt after learning about the cost incurred. Data analysis was carried out using IBM's SPSS, and missing data were deleted list-wise.

Results

Table 4 shows mean initial elevation and mean change in elevation for participants in each of the three cost conditions, and Table 5 shows the means for male and female participants, and religious and non-religious participants.

Table 4

Initial Moral Elevation and Change in Moral Elevation by Condition

Variable	Condition					
	High		Moderate		Low	
	<i>n</i>	<i>M(SD)</i>	<i>n</i>	<i>M(SD)</i>	<i>n</i>	<i>M(SD)</i>
Initial	38	5.42 (1.35)	34	5.26 (1.21)	54	5.50 (1.19)
Change	38	5.53 (1.50)	34	4.91 (1.38)	54	3.90 (1.45)

As with study 1, incompleteness of data prevented ordinal analyses, and so elevation scores were treated as scale variables. Because neither initial elevation nor change in elevation were normally distributed (respectively, $D(130) = .13$, $p < .001$, and $D(130) = .19$, $p < .001$), non-parametric analyses were adopted. Bonferroni corrections were applied to compensate for multiple comparisons.

Two Kruskal-Wallis tests were performed to establish whether initial moral elevation or change in moral elevation varied significantly across the three conditions (it was anticipated that initial elevation would not vary, and it was hypothesized that change in elevation would). Initial moral elevation did not vary significantly across the three cost conditions ($H(2) = .90$, $p = .64$), while change in elevation did vary significantly ($H(2) = 25.62$, $p < .001$). Follow-up Mann-Whitney tests revealed that participants in the high and medium cost conditions reported significantly larger increases in elevation than participants in the low cost condition (respectively, $U = 38.15$, $p < .001$, and $U = 23.45$, $p = .003$), although the difference in change in elevation between high and medium cost condition participants was non-significant – $U = 14.71$, $p = .09$. The mean differences are illustrated in Figure 3.

There was no significant difference between male and female participants for either initial elevation ($U = 2307$, $p = .35$) or change in elevation ($U = 2488$, $p = .07$). Similarly, there was no significant difference in either initial elevation or change in elevation between religious and non-religious participants; respectively, $U = 1971.5$,

$p = .37$, and $U = 1979$, $p = .34$. Age was also non-normally distributed; two Spearman's rank-order correlations revealed a significant positive relationship between age and initial elevation ($r_s(129) = .27$, $p = .002$) but no significant relationship between age and change in elevation ($r_s(129) = .06$, $p = .50$).

Table 5

Initial Moral Elevation and Change in Moral Elevation by Sex and Religious Affiliation

Variable	Sex				Religious			
	Female		Male		Yes		No	
	<i>n</i>	<i>M(SD)</i>	<i>n</i>	<i>M(SD)</i>	<i>n</i>	<i>M(SD)</i>	<i>n</i>	<i>M(SD)</i>
Initial	63	5.60 (1.21)	67	5.24 (1.25)	80	5.49 (1.25)	45	5.29 (1.20)
Change	63	4.79 (1.53)	67	4.49 (1.66)	80	4.73 (1.57)	45	4.44 (1.63)

Figure 3. Change in Moral Elevation by Cost Condition

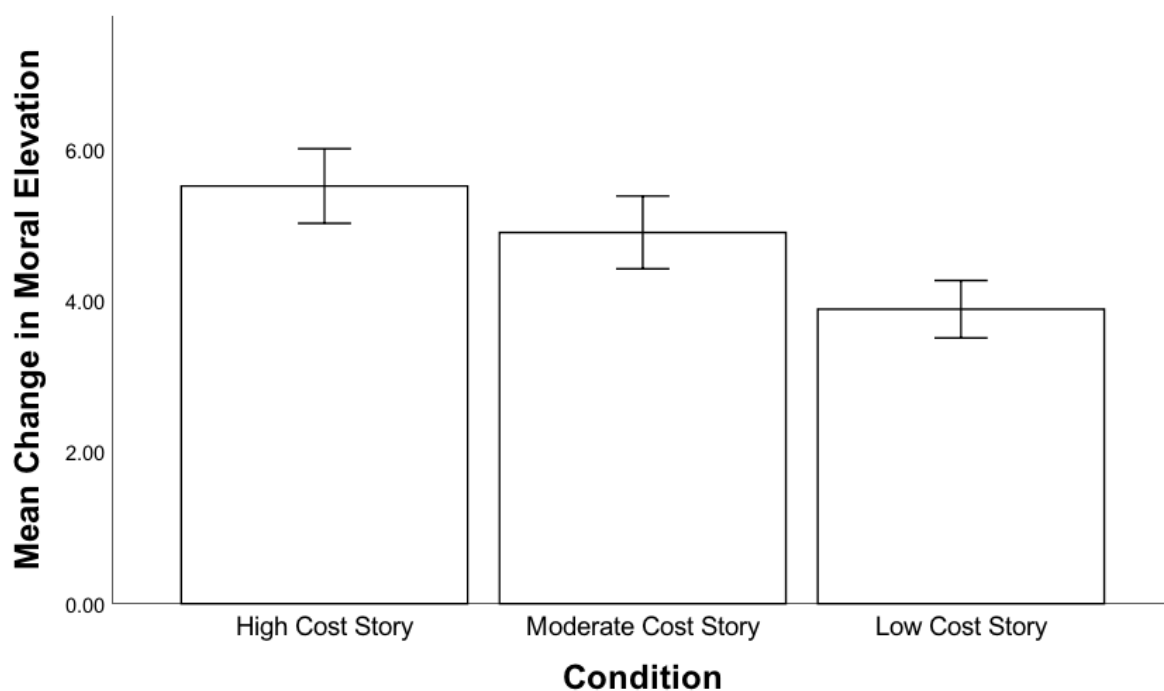


Figure 3. Mean change in elevation score for participants in each story condition. Error bars represent 95% confidence intervals.

Discussion

The results of study 2 offered good support for prediction 2; participants reported a greater positive change in their experiences of moral elevation when they learned that a benevolent act had cost the actor more dearly. There was a clear linear relationship between cost incurred by the actor and subsequent change in moral elevation. Participants in the low cost condition reported a slight decrease in moral elevation, and participants in the high cost condition reported a higher increase in elevation than participants in the moderate cost condition. These group differences were not present for the pre-manipulation, initial reports of moral elevation, giving confidence that the differences in change in moral elevation represent an experimental affect. It seems that moral elevation may indeed be sensitive to features of a benevolent act which indicate that the actor values the welfare of the recipient highly, suggesting that it is functionally integrated with our welfare trade-off ratio estimator. This apparent synergy between moral elevation and another mechanism in our social-cognitive architecture supports the idea that elevation plays a functional role in regulating our other-regarding behaviours.

As in study 1, female participants and participants who reported religious affiliations expressed stronger feelings of moral elevation than male or non-religious participants. Although only the higher initial elevation of women was statistically significant, the means summarized in Table 4 reflect what seem to be steady trends. Unlike in study 1, there was a significant positive correlation between age and participants' initial reported moral elevation.

Conclusion

Studies 1 and 2 aimed to test predictions about which particular aspects of others' behaviour should be expected to elicit moral elevation, or to elicit it more strongly, if the emotion is indeed a cognitive mechanisms sculpted by natural selection to respond to indications of high partner-value in others.

It was reasoned in Chapter 2 that individuals who tend to directly confer benefits on others make better (more profitable) social-relationship partners than individuals who tend to exhibit norm or rule following behaviour (all else equal). On that basis, it was expected that moral elevation should be more sensitive to – i.e. more easily elicited by – others' benefit conferring behaviours than by their simple norm-following behaviours. The results of study 1 failed to support this prediction,

but study 3B (reported in the next chapter) will offer another chance to test prediction 1 using a less problematic control vignette.

It was also reasoned in Chapter 2 that moral elevation, if it is a functionally integrated part of our social-cognitive machinery, should be sensitive to information about how much an actor values the welfare of the recipient in a benevolent act. So-called welfare trade-off information (i.e., information about how much of her own welfare an actor is willing to sacrifice in order to increase the welfare of the recipient) is known to be something we track and respond to in evaluating others' worth as potential social-relationship partners (Delton et al., 2012; Delton & Robertson, 2012), and so it was predicted that there would be a positive relationship between the proportionate cost an actor incurs in performing a benevolent act and the degree of moral elevation the act elicited. This prediction was supported by the results of study 2. A further test of prediction 2 will be offered in study 4, reported in the next chapter.

Having offered preliminary investigations of the types of behaviours which elicit moral elevation, Chapter 4 will turn to investigations centered around which types of actors and recipients moral elevation may be particularly sensitive to.

Chapter 4: Empirical Tests of Predictions 3 and 4

Abstract

Chapter 4 reports the results of three studies. Studies 3A and 3B aimed to test the prediction that beneficiary-specific incentives for a benevolent actor's behaviour should decrease the power of the behaviour to elicit moral elevation. Study 4 aimed to test the prediction that benevolent behaviour will be more morally elevating when it is performed by actors who are otherwise relatively socially attractive. Neither Study 3A nor study 3B found support for the prediction – implications are discussed. Study 4 found some support for the prediction – study limitations are discussed.

Studies 3A and 3B: Is Moral Elevation Sensitive To The Presence Of Incentives For The Benefactor's Behaviour?

Study 3A

Study 3A aimed to test prediction 3:

Prediction 3: A benevolent act will elicit stronger feelings of moral elevation if there are no obvious incentives for the actor to value the welfare of the recipient

Design

In contrast to studies 1 and 2, study 3A used a within-subjects design. All participants were asked to read two short stories and answer a series of questions intended to gauge which story they found more morally elevating. One story described someone performing a benevolent act in a situation where there was a clear incentive for the actor to value the recipient's welfare, and one story described a benevolent act in a situation with no clear incentive. In line with prediction 3, it was anticipated that participants would report finding the story in which there was no clear incentive for the actor's behaviour more morally elevating than the alternative story.

As in study 1, participants in study 3A were asked to complete a 9-item forced choice questionnaire measure intended to test prediction 6 (again, to maximize data-collection efficiency). The results from that measure will be presented in Chapter 6, along with a suitable re-cap of sample characteristics and data processing methods.

Method

Participants

106 participants were recruited via Amazon's MTurk between 24-25/02/2019, and were paid \$1 for their time. Ethical approval was granted for this study on 23/01/2019.

All participants (43% female, age range 22-69, $M = 36.44$, $SD = 10.64$) were from the USA, and 54% indicated that they held no religious affiliation. Participant attention was checked at three points in the survey; two comprehension-based questions following the story texts they were asked to read, and one free text

question at the end of the survey asking participants to answer honestly whether they had been paying attention throughout the survey. Participants who gave incorrect answers for either of the comprehension questions, and those who gave anything other than an affirmative answer to the free text question were excluded from analyses. The remaining N was 102.

Measures

Two short vignettes describing benevolent acts were used for this study. In story 1, Sam saves Ian's valuables when they fall into the river on a company team building day. In story 2, Adam goes out of his way to give a man a lift home on a rainy night. In order to avoid confounds produced by any differences between these two scenarios, each story was written in two versions – in one version (A), the recipient had no qualities that would obviously incentivize the actor to value his welfare, and in the other version (B) the recipient did have a quality that would obviously incentivize the actor to value his welfare. In story 1A, Ian is described as Sam's boss, and in story 1B Ian is described as a very junior colleague of Sam. In story 2A, the man whom Adam helps is described as a well-heelled old friend from college, and in story 2B the man is described as a stranger who was down on his luck (each version is given in full in Appendix B).

The scale used to measure which story participants were more morally elevated by was based around Vianello et al.'s (2010) 8-item elevation scale, described in the measures section of study 1 in Chapter 1 (p. 30) above. Participants were asked, for each of the 8 sensations in the moral elevation scale, which story, if either, had elicited it most strongly. For example:

1) Did either story make you feel more open towards others? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, the first story
- b) Yes, the second
- c) No, neither story

The full measure is included in Appendix B. Participants were also asked to provide basic demographic information.

Procedure

Participants were asked to follow a hyperlink from MTurk to a survey hosted on Qualtrics. They were randomly assigned, by Qualtrics, to one of two conditions: in condition 1, participants were asked to read story 1A and story 2B, and in condition 2 participants were asked to read story 1B and story 2A. Following the stories, and the attention check comprehension questions, participants were asked to complete the comparative moral elevation measure.

The comparative moral elevation measure was scored by giving participants 1 point each time they chose an answer that equated to being more elevated by the story in which the benefactor had an incentive ('A' version stories), -1 point each time they chose an answer that equated to being more elevated by the story in which the benefactor had no incentive ('B' version stories), and 0 point each time they chose the 'neither story' option. Scores could be anywhere between -8 and 8. A score < 0 is taken to indicate that the participant was more elevated by the B story (where there was no incentive), and a score > 0 is taken to indicate that the participant was more elevated by the A story (where there was an incentive). A score of 0 is taken to mean that the participant was no more elevated by one story than by the other.

Results

Did participants find one story more elevating than the other?

In total, 45.1% of participants had scores < 0 , 44.1% had scores > 0 , and 10.8% had a score of 0, meaning that marginally more participants reported finding the B story (where the protagonist had no obvious incentive for his actions) to be more elevating than the A story (where the protagonist had an incentive to help the recipient). The magnitude of this difference, however, was trivial. Table 6 is the contingency table showing which story was found to be more morally elevating by participants in each condition, and Figure 4 shows the number of participants who chose the incentive and no-incentive stories across both conditions.

The data from participants in condition 1 conform to the expectations of prediction 3 – participants were overwhelmingly more likely to rate the story in which the protagonist had no clear incentive for his actions as more morally elevating than the story in which there was an obvious reason for the protagonist to value the welfare of the recipient (which in this condition was story 2B). The answers of

participants in condition 2, however, ran in the opposite direction – these participants were far more likely to rate the story with the clear incentive as more elevating than the story with no clear incentive (which in this condition was story 2A). A chi-square test confirmed that there was a significant association between condition and choice of which story was more elevating; $\chi^2 (2) = 31.54, p < .001$. In short, participants in both conditions were significantly more likely to rate story 2 as more morally elevating than story 1.

Table 6
Story Chosen as More Elevating by Condition

Condition	A story <i>n</i> (%)	B Story <i>n</i> (%)	Neither <i>n</i> (%)
Condition 1	7 (14.9)	34 (72.3)	6 (12.8)
Condition 2	38 (69.1)	12 (21.8)	5 (9.1)
Total	37 (44.1)	46 (45.1)	11 (10.8)

Were there any age, sex, or religiosity differences in moral elevation?

In order to investigate whether there were any relationships between moral elevation and sex, age, and religiosity, it was necessary to score the comparative elevation measure using a different method. As described above, the comparative moral elevation measure was initially scored by giving participants 1 point each time they chose the A story as more elevating, -1 point each time they chose the B story as more elevating, and 0 points each time they picked ‘neither’. This left participants with scores between -8 and 8, and allowed for the categorical comparisons necessary to test prediction 3; a score > 0 was taken to indicate that a participant found the A story more elevating, a score < 0 that they found the B story more elevating, and a score of 0 that they found neither story more elevating than the other. Looking for sex, age, and religiosity differences, however, requires a continuous dependent variable.

To gauge how elevated participants felt overall (in response to both stories combined), they were given a score of 1 each time they answered that either the A story or the B story had evoked an elevating sensation, and a score of 0 each time

they picked 'neither'. This left participants with a moral elevation score between 0 and 8, with higher values taken to indicate that the participant experienced stronger feelings of moral elevation. Mean moral elevation scores by sex and religious affiliation are shown in Table 7.

Figure 4. Story Choice




Figure 4. Total number of participants who chose the incentive-present and incentive-absent stories as more morally elevating.

Table 7

Moral Elevation by Sex and Religious Affiliation

Group	<i>n</i>	<i>M(SD)</i>	<i>t(df)</i>
Sex			
Male	57	5.31 (2.43)	
Female	44	6.05 (2.43)	1.72 (98.95)
Religious			
Yes	48	6.38 (1.47)	
No	54	4.87 (2.61)	3.64 (85.31)*
Total	102	5.58 (2.27)	

Note: * $p < .05$

Once again, female participants expressed stronger feelings of moral elevation than male participants, although the difference only approached significance; $t(98.95) = 1.72$, $p = .08$, Cohen's $d = .30$ (equal variances not assumed). Religious participants also expressed significantly stronger feelings of moral elevation than non-religious participants; $t(85.31) = 3.64$, $p < .001$, Cohen's $d = .71$ (equal variances not assumed). There was no significant correlation between age (which again was non-normally distributed) and moral elevation; $r_s(102) = .11$, $p = .26$.

Discussion

Contra prediction 3A, participants in study 3 were not more likely to express feelings of moral elevation in response to the B versions of the stories than they were in response to the A story versions. In fact, participants were overwhelmingly more likely to report that story 2 was more elevating than story 1, regardless of which variants they read, suggesting that the manipulation may have failed. Best efforts were made to match the two vignettes as closely as possible, but there may still have been too many confounds to produce adequately interpretable results. If there was any effect of the presence or absence of incentives for the protagonist to value the beneficiary's welfare, it seems it would have been drowned out by the effects of other elements of the two stories. It should also be considered that participants may have been expressing a recency bias in their answers; stories 2A and 2B were the second stories presented in each condition. The choice to use a within-subjects design was motivated by cost efficiency, but the drawbacks of repeated measures

could have negatively affected the power of this design to detect the sought-after effect.

Once again, however, there was evidence contributing to the pattern of a stable sex-difference in moral elevation, with female participants scoring marginally-significantly higher than male participants. As in study 1, religious participants reported significantly stronger feelings of moral elevation than non-religious participants. It bears noting, though, that the moral elevation scale used was less than optimal for these analysis. In retrospect it would have been preferable to have given participants a wider range of response options, so as to be able to more accurately quantify the strength of their feelings of elevation. As things stood, with answers to each of the moral elevation scale items being effectively binary, it was only possible to garner a relatively crude idea of how elevated participants felt.

Study 3B

Design

Study 3B also aimed to test prediction 3, but unlike study 3A, 3B employed a between-subjects design. Participants were asked to rate how morally elevated they felt after reading stories about benevolent acts which occurred between family members, friends, or strangers. In line with prediction 3, it was expected that participants who read the stories about benevolence between strangers would report stronger feelings of moral elevation than participants who read about benevolence between family members or friends.

Study 3B also afforded an opportunity to further test prediction 1:

Prediction 1: Moral elevation is more likely to be elicited by others' benevolent behaviour than by behaviours which are morally praiseworthy but non-benevolent

Participants were asked to read three short stories in total, two describing benefit conferring behaviours, and one describing behaviour which would be considered morally laudable but which did not involve benefitting another individual. After each story, they were asked to indicate which emotional state they felt the story had evoked in them. In line with prediction 1, it was anticipated that participants

would only choose moral elevation as their response to the two benevolent stories, and not to the morally laudable but non-benevolent story.

Method

Participants

163 participants were recruited through Crowd Flower between 25-29/11/2014, and were paid \$0.50 for their time. Ethics approval was granted for the study on 09/09/2014. 64% of participants were female, and the age range was 18 – 69 ($M = 37.43$, $SD = 12.48$). Participant nationality and religious affiliation were not recorded for study 3B.

Measures

The three stories presented to participants were based on real news stories, and the only details changed were the relationships between the protagonists and the recipients.

One story acted as a control was not varied between participants. This described a young lady continuing her passion for and improvement in CrossFit (exercise) training, following a car accident after which her arm was amputated. The expectation was that this story would elicit inspiration, and not moral elevation, despite describing an individual with a number of morally positive virtues (e.g., strength of character, determination, optimism).

The other two stories were found by searching the internet for news items tagged as 'restoring faith in humanity'. They describe noteworthy acts of helpfulness or kindness, and were expected to elicit moral elevation. The first ('elevating story 1') described a young boy who wrote a book and sold copies of it to raise money for research into a rare disease that one of his friends suffers from. The second story ('elevating story 2') described a marine who carried a young para-athlete across the finish line of a marathon after the young lad's prosthetic leg failed. Both elevating stories were varied three ways: in condition 1, the 'high incentive' condition, the benefactor and recipient were described as brothers; in condition 2, the 'moderate incentive' condition, the benefactor and recipient were described as friends; and in condition 3, the 'low incentive' condition, the benefactor and recipient were described

as strangers. All three versions of both elevating stories, and the story used as a control, are given in full in Appendix C.

Moral elevation was measured using Vianello et al.'s (2010) 8 item elevation scale. Participants were asked to rate, using a 5-point Likert scale (1 = 'not at all', 5 = 'very much'), how strongly they felt each of the 8 sensations described in the scale. Cronbach's α was .91.

Finally, participants were given the opportunity to donate their participation fee to a charity (Medecins Sans Frontieres); this was included as a behavioural measure of moral elevation, following other researchers who have taken charitable donation as indicative of elevation's pro-social motivational output (Freeman et al., 2009) (Aquino et al., 2011) (Thomson & Siegel, 2013).

Procedure

Participants were directed, via hyperlink, to a questionnaire hosted on Survey Monkey. The questionnaire was described to participants as a study investigating emotional reactions to news stories. In order to reduce the chance of producing demand characteristics, participants were told that the researcher was interested in three emotions (or 'feelings'); having one's faith in humanity restored, being inspired to reach one's full potential, and being amused.

'Having your faith in humanity restored' was used in place of 'feeling morally elevated' because moral elevation is not a commonly used phrase, and it is one with a somewhat checkered history – the term was more popular in the second half of the 19th century, when it was used to refer to a deliberate attempt at bettering, or morally educating, others who were deemed 'less fortunate' (often women, ethnic minorities, and the poor). The term 'faith in humanity restored', on the other hand, has become a popular internet meme over the past 10-15 years. Used to describe the reaction to an action or event which improves the observer's opinion of people in general, the phrase fits well with descriptions given in the elevation literature of the appraisals which accompany experiences of moral elevation. It is an apt synonym for having one's "views about humanity changed in a more optimistic way" (Haidt, 2000), for feeling "optimistic about humanity" (Haidt, 2003a), and for realizing that "there are good people in the world, there is good in people" (Haidt, 2003a).

The other two emotions, being inspired and being amused, were included as decoys, to mask the research intention of the study. The use of these two decoys in particular was motivated by the similarity between these two emotion states and moral elevation; all three are positive, pleasant affective states, evoked by witnessing others' actions (Algoe & Haidt, 2009). The decoy emotions could therefore stand as partial manipulation checks, able to offer indications of whether the elevating materials used were eliciting moral elevation, or some other type of positive, other-orientated reaction.

A further benefit of the inclusion of 'being inspired' as a decoy was the opportunity to replicate the differentiation of inspiration and elevation that Algoe and Haidt (2009) achieved in their work exploring emotional reactions to excellence in others. They found evidence to support characterizing moral elevation and inspiration (which they conceive of as the motivational output of 'admiration') as reactions to distinct types of social stimuli. Whereas moral elevation is brought about by "others exceeding standards of virtue", admiration (and its concomitant motivational effect, inspiration) is brought about by "others exceeding standards of skill or talent" (p.108). With evidence grounding fine-grained distinctions between positive emotions still something of a lacuna in the literature (Strohlinger et al., 2011), opportunistic verifications of the distinctions which have so far been drawn can help to firm our confidence in the coherence of the phenomena being explored.

Participants were given a description of each emotion, and asked to read them carefully as "different people may have slightly different understandings of words related to feelings or emotions" (participant instructions). As a further precaution against demand characteristics, the descriptions did not make reference to stimuli that would normally elicit the feelings described, but instead detailed the sensations that accompany the emotion and the motivations that typically follow from experiencing it. The full descriptions ran as follows:

Please read the following descriptions of 3 emotions. Each description is broken into two parts: what sensations or feelings the emotion produces, and what the emotion typically makes us want to do.

1. Having your faith in humanity restored

Feeling: Optimistic and positive about humanity or human nature. Warm, uplifted, open, moved, touched.

Wanting to: Show others that human nature is good by setting an example. Being moved to do positive, caring, kind, or helpful things for others.

2. Amusement or mirth

Feeling: Light-hearted, high-spirited, entertained.

Wanting to: Laugh, play, engage in merriment, do something fun.

3. Being inspired to reach your full potential

Feeling: Energized, motivated, and positive about the possibility of achieving goals.

Wanting to: Take on new challenges, work harder to reach targets, or throw yourself into activities that matter to you.

To check attention and understanding, participants were asked to indicate for each emotion whether they thought it was positive/pleasant, negative/unpleasant, or were unsure about its valence.

Following this, participants were presented with three news stories (sequentially). Each participant saw the (putatively) inspiring story, and was randomly assigned to see one version of each of the (putatively) elevating stories; the order of viewing was fixed as elevating story 1, inspiring story, elevating story 2.

After each story, they were asked to indicate which of the three previously described emotions they had experienced most strongly while reading (with 'this story did not make me feel any of the above' given as a fourth response option). Choosing one of the three emotions as their response took participants to a page with further probes about the strength and quality of the emotional experience. The probes for having one's faith in humanity restored comprised the 8 item moral elevation scale from (Vianello et al., 2010). The probes for amusement and inspiration were included only to authenticate the deception that the study was aimed at exploring more than one emotion; there were 7 amusement probes, and 8 inspiration probes, to give face-validity in juxtaposition with the elevation scale. Participants who chose the option 'this story did not make me feel any of the above' after reading the story were directed straight to this page.

Participants' moral elevation scores were computed by summing their responses on the 5-point Likert scale to the 8 moral elevation scale items, and dividing that sum by 8. This gave each participant a moral elevation score between 1 and 5, with higher scores indicating stronger feelings of moral elevation. Each participant completed the moral elevation scale a maximum of 3 times, because they were given the option to choose 'this story restores my faith in humanity' after reading each story. Only the scales which were completed after participants read the two benevolent stories (stories 1 and 2) were computed and used for the analyses below. Data analysis was conducted using IMB's SPSS package.

Results

Did the three stories elicit the anticipated emotional responses?

In line with prediction 1, the two putatively elevating stories were overwhelmingly described by participants as such, and the putatively inspiring story was equally overwhelmingly described as inspiring rather than elevating. Table 8 and Figure 5 show the choice frequency for each response option across all three stories; story 1 (the boy who writes the book), story 2 (the marine who carries the boy), and story 3, the control story (the amputee who succeeds at Cross Fit).

Were there any age or sex differences in moral elevation?

In order to perform these analyses, participants' moral elevation scores for each of the elevating stories were combined and divided by 2, to produce a total elevation score for each participant. These scores were normally distributed; $D(79) = .07, p = .20$. In contrast to study 1, there was no significant difference between the mean elevation scores of male and female participants (male $M = 3.34, SD = .59$, female $M = 3.49, SD = .91$); $t(76) = -.75, p = .46$. Neither was there any significant correlation between age and moral elevation; age was once again non-normally distributed, but Spearman's rank-order correlation was $r_s(79) = .01, p = .90$.

Was there a relationship between story condition and moral elevation?

The mean moral elevation scores for participants in each condition for the two morally elevating stories are given in Table 9. Contra prediction 3, two one-way ANOVAs found no significant main effect of condition on moral elevation for either

story: for elevating story 1, $F(2,99) = 1.11$, $p = .33$; and for elevating story 2, $F(2,109) = .26$, $p = .77$.

Table 8
Response Option Choices for Each Story

Story	Response Option			
	Elevating <i>n</i> (%)	Inspiring <i>n</i> (%)	Amusing <i>n</i> (%)	None <i>n</i> (%)
Story 1	107 (65.6%)	37 (22.7%)	5 (3.1%)	14 (8.6%)
Story 2	116 (71.2%)	20 (12.3%)	6 (3.7%)	21 (12.9%)
Control	7 (4.3%)	102 (62.6%)	10 (6.1%)	44 (27%)

Was there a relationship between moral elevation and participant donations to charity?

In total, only 19 participants chose to donate their participation fee to MSF. There was no significant difference between the mean elevation scores of participants who chose to donate and those who chose not to (donation $M = 3.30$, $SD = .83$, no donation $M = 3.47$, $SD = .90$); $t(76) = -.61$, $p = .54$.

Table 9
Moral Elevation by Condition

Story	<i>n</i>	Strangers	<i>n</i>	Friends	<i>n</i>	Brothers	<i>F</i> (<i>df</i>)	<i>p</i>
		<i>M</i> (<i>SD</i>)		<i>M</i> (<i>SD</i>)		<i>M</i> (<i>SD</i>)		
Story 1	26	3.42 (.76)	46	3.28 (.92)	30	3.07 (.96)	1.11 (2, 99)	.33
Story 2	30	3.54 (.79)	44	3.61 (.93)	38	3.37 (.98)	.26 (2, 109)	.77

Figure 5. Emotional Response to Each Story




Figure 5. Number of participants reporting moral elevation, amusement, or inspiration to each of the three news stories.

Discussion

Study 3B rectified the major methodological flaws that were present in study 3A – it used a between-subjects design, allowing the same vignette to be used across conditions with only very minimal changes, and the absence of repeated measures also removed any possibility of a presentation-order effect – but there was still no support found for prediction 3. There were no significant differences between the elevation scores of participants in the different story conditions, lending good support to the counter-hypothesis, that moral elevation is not sensitive to the presence or absence of motivations for a benevolent actor to value the welfare of the

recipient. If it is in fact a cognitive adaptation, moral elevation would seem not to have the level of functional sophistication assumed by prediction 3.

Study 3B did, however, provide support for prediction 1. Although study 1 failed to find support for the idea that moral elevation is elicited exclusively by benefit-conferring, and not by non-benevolent but otherwise morally laudable, behaviours, study 3B used a less ambiguous vignette as its control and found that participants were vastly less likely to express feelings of moral elevation in response to others' moral-but-non-benevolent behaviour. Only 4.3% of participants said that the story about the amputee whose extraordinary perseverance and positivity led her to break her pre-amputation Cross Fit records was morally elevating, compared with the 71.2% and 65.6% of participants who said that the stories about the marine helping the young triathlete and the boy writing and selling a book to help his friend (respectively) evoked elevation.

This study did not find a significant sex-difference in reported moral elevation. Although the means trended in the same direction as those in studies 1, 2, and 3A, the difference was only very slight and was far from statistically significant. As the result of an unfortunate oversight, participant religiosity was not recorded, and so it was not possible to look for any difference in reported elevation between religious and non-religious participants.

Study 4: Is Moral Elevation Sensitive To Socially-Attractive Characteristics Of The Actor?

Study 4 aimed to test prediction 4, and it provided an opportunity to further test prediction 2.

Prediction 4: A benevolent act will elicit stronger feelings of moral elevation when the actor has other traits indicative of high partner-value

Prediction 2: A benevolent act will elicit stronger feelings of moral elevation if the actor has to pay a higher proportionate cost to perform the act

Design

Study 4 employed a 2 x 3 between participant experimental design. Participants were asked to rate their feelings of moral elevation after reading a short story describing a benevolent act, accompanied by a picture of the protagonist's face; the act itself was either described as low cost or high cost, and the face accompanying the story was either a high, moderate, or low dominance male. Socially dominant males are physically formidable individuals, capable (at least ancestrally) of acquiring and holding resources, and likely to emit a number of positive externalities which would benefit their social allies (Jensen & Petersen, 2011; Sell et al., 2009); higher dominance males, then, are more attractive social partners than lower dominance males, *ceteris paribus*.

In line with prediction 4, it was anticipated that participants who read the story which was accompanied by the high-dominance face would report the strongest feelings of moral elevation, and participants who read the story accompanied by the low-dominance face would report the weakest feelings of elevation. In line with prediction 2, it was also anticipated that participants who read the high-cost version of the story would report stronger feelings of elevation than participants who read the low-cost version.

Method

Participants

Participants ($N = 315$, 56.2% female) were recruited through Crowd Flower and Amazon MTurk, between 15-17/12/2013, and were paid \$0.35 for taking part in the study. Ethical approval was granted for the study on 30/10/2013.

Participant age was recorded in ordinal blocks (e.g., 18-25, 25-35, etc.); 35.9% were aged between 25 and 35, 20% were between 35 and 45, 18.4% between 18 and 25, and 15.9% between 45 and 55. 66.5% of participants indicated that they held some kind of religious affiliation. The majority of participants were from the USA (72.1%), 20% were Canadian, 5.4% from the UK, and the remaining 2.5% from a handful of other countries.

Measures

Moral elevation was elicited using a vignette about a young man helping an elderly man to clear his driveway of snow. Participants saw either a low cost version of the story, in which the young man has to invest a moderate amount of effort to help the elderly man, or a high cost version, in which the young man has to invest substantially more time and energy. The first paragraph was identical for both story versions, and ran as follows:

Both Story Conditions

Alex was being given a lift home by his friend Martin one winter afternoon after the two had met for lunch together in town. As they drove they passed an elderly man, his coat collar turned up against the cold wind that was picking up, shovelling snow from his driveway. Alex tapped Martin's shoulder and said, 'You can just let me out here, thanks Martin'. They were about half a mile from Alex's house, but Martin assumed that Alex wanted to stop in at his girlfriend's place, which was close-by; he pulled over and let Alex out. As Martin drove off, Alex jogged back to the man with the shovel to offer help. Alex and the elderly man chatted while they took it in turns to shovel the snow – Alex making sure he himself did the majority of the work. Alex learned that the man had wanted to clear the snow because his daughter was driving down to visit that evening with his new grandchild.

The stories varied in their closing lines:

Low cost story

The snow-fall had been quite light, around 4 inches, and the wind eased as they worked, making the job easier. In total, Alex spent 15 minutes helping the man clear the 5 yards of driveway. When the job was finished, Alex said goodbye, the man thanked him for his kindness, and Alex walked home.

High cost story

[Identical to the low cost story until the final four lines]:

The snow-fall had been heavy, around 8 inches, and the worsening wind made the job increasingly difficult. About half way up the 30 yard track, Alex grew concerned about the elderly man being out in the cold and the wind, so he convinced him to go

inside the house and have a warm drink, while Alex remained outside to finish the job. In total, Alex spent two and a half hours clearing the snow. When the job was finished, Alex said goodbye, the man thanked him for his kindness, and Alex walked home.

Perceptions of the protagonist's dominance were manipulated using pictures of male faces accompanying the stories and labeled with the protagonist's name. These pictures (shown in Figure 6) were taken from a database of 25 maximally distinct male faces, created using FaceGen modeler software (version 3.1) (Todorov, Dotsch, Porter, Oosterhof, & Falvello, 2013). Each identity in the database has been manipulated on the dimension of dominance (Todorov & Oosterhof, 2011) to create 7 versions (-3SD, the lowest dominance, to +3SD, the highest dominance). One identity was chosen (number 25) to represent Alex; the -1SD version served as the low dominance stimulus, the +1SD version was the moderate dominance stimulus, and the +3SD version provided the high dominance stimulus.

Moral elevation was measured using the 8 item moral elevation scale from Vianello et al. (2010); Cronbach's α was .86. Participants were also asked to answer a 7 item questionnaire intended to gauge their perceptions of Alex's value as a social partner. The measure included items such as 'I would like to spend time socially with Alex', and 'Alex would make a good friend', scored on a 7-point Likert scale (full questionnaire included in Appendix D); Cronbach's α for this scale was .72.

Figure 6: Picture Stimuli Used For High, Moderate, and Low Dominance Conditions




Figure 6: Left to right, 'Low Dominance' (-1SD), 'Moderate Dominance' (+1SD), and 'High Dominance' (+3SD)

Procedure

Participants were directed via hyperlink to a questionnaire hosted on Survey Monkey. Participants were randomly assigned by Survey Monkey to one of six conditions. Each participant was asked to read either the low cost or the high cost version of the altruistic vignette. Alongside the story, participants were presented with a picture of a male face which was either high dominance, moderate dominance, or low dominance. After reading the story, participants were asked to rate their perceptions of the protagonist's dominance, their experiences of moral elevation, and their perceptions of the protagonist as a prospective social partner.

Data analysis was carried out using IBM's SPSS. There was a relatively high proportion of missing data, and these were deleted list-wise in all analyses.

Results

Participants completed an 8-item moral elevation measure, scored on a Likert-scale running from 0-5; summing and averaging these scores gave each participant a moral elevation score between 0 and 5, with higher values indicating that participants experienced stronger feelings of moral elevation. Mean moral elevation scores of participants in each picture and story condition, as well as male and female, and religious and non-religious participants, are given in Table 11. Moral elevation scores were somewhat negatively skewed, creating a significantly non-normal distribution; $D(296) = .06$, $p = .03$. Reverse Log transformation did not improve the normality test-results, but the absolute value of skew was relatively small (Z -skew = -1.70), and it is possible that these tests were too sensitive for the relatively large study N .

Dominance manipulation check

In order to establish whether the picture stimulus each participant was given affected their perceptions of Alex's dominance in the expected fashion, participants were asked to rate (using a 5 point Likert scale) how likely they thought Alex would be to win in a physical confrontation with a man of a similar age. Mean dominance ratings for male and female participants and participants in each picture and story condition are given in Table 10.

Table 10

Dominance Ratings by Sex, Picture Condition, and Story Condition

Group	<i>n</i>	<i>M(SD)</i>	<i>t(df)</i>
Picture Condition			
Low	91	3.27 (.72)	
High + Moderate	223	3.58 (.78)	3.15 (311)*
Story Condition			
High	166	3.58 (.88)	
Low	148	3.39 (.73)	2.14 (312)*
Sex			
Male	137	3.43 (.78)	
Female	177	3.54 (.77)	1.20 (312)

Note: * $p < .05$

A one-way ANOVA found a significant effect of picture condition on ratings of Alex's dominance; $F(2, 311) = 5.25, p = .006$. Planned contrasts showed a significant difference between the means of the high and moderate dominance picture conditions and the low dominance picture condition ($t(311) = 3.15, p = .002$), but no significant difference between the high dominance and moderate dominance picture conditions ($t(311) = -.58, p = .56$). These results suggest that the manipulation of Alex's dominance was only partially successful. Consequently, participants from the high and moderate dominance picture conditions were combined into one group (displayed in Table 9 as 'High + Moderate'). All of the following analyses contrast the means of this combined group with the means of the participants in the low dominance picture condition – referred to hereafter as the high dominance picture condition and the low dominance picture condition respectively.

Unexpectedly, there was a significant difference in ratings of Alex's dominance between the high and low cost story conditions, with participants in the high cost condition rating Alex as significantly more dominant than participants in the low cost condition; $t(312) = 2.14, p = .03$. A two-way ANOVA was conducted to check whether there were any interaction effects between picture condition and cost condition (no interaction effects were expected); as anticipated, there were significant main effects for both picture and story condition (respectively, $F(2, 308) = 5.54, p = .004$, and $F(1, 308) = 5.23, p = .02$), but the interaction effect was non-significant; $F(2, 308) = .37, p = .70$.

There was no significant difference in dominance ratings between male and female participants. Although women rated Alex as more dominant than men did, the difference was non-significant; $t(312) = 1.20, p = .23$.

Table 11
Mean Moral Elevation by Picture Condition, Story Condition, Sex, and Religiosity

	Group	<i>n</i>	<i>M(SD)</i>
Picture Condition	High	213	3.31 (.79)
	Low	84	3.24 (.73)
Story Condition	High	154	3.33 (.79)
	Low	143	3.24 (.75)
Sex	Male	131	3.18 (.72)
	Female	166	3.37 (.80)
Religious	Yes	192	3.39 (.77)
	No	103	3.11 (.74)

Did story condition, picture condition, sex, or religiosity predict participants' moral elevation?

Mean elevation scores are given in Table 11. A backwards stepwise, linear multiple regression was conducted with cost condition, dominance condition, sex, and religiosity as predictors, and moral elevation as the outcome variable. Table 12 shows the resulting coefficients. Sex was the only significant predictor in the model ($F(1,293) = 4.58, p = .03$), although it only accounted for 2% of the variance in moral elevation ($R^2 = .02$).

Table 12

Regression Coefficients: Cost Condition, Dominance Condition, Sex, and Religiosity as Predictors of Moral Elevation

Predictors	B	SE B	β
Step 1			
Constant	3.26	.15	
Cost: Low vs. High	.08	.09	.05
Dominance: Low vs. High	.09	.10	.06
Sex: Female vs. Male	-.19	.09	-.12*
Religiosity: No vs. Yes	.001	.13	<.001
Step 2			
Constant	3.27	.11	
Cost Condition: Low vs. High	.08	.09	.05
Dominance Condition: Low vs. High	.09	.10	.06
Sex: Female vs. Male	-.19	.09	-.12*
Step 3			
Constant	3.31	.09	
Dominance Condition: Low vs. High	.09	.10	.05
Sex: Female vs. Male	-.20	.09	-.13*
Step 4			
Constant	3.37	.06	
Sex: Female vs. Male	-.19	.09	-.12*

Note: R^2 for step 1 = .02 ($p = .19$), ΔR^2 for step 2 = <.001, ($p = .99$), ΔR^2 for step 3 = .003 ($p = .38$), ΔR^2 for step 4 = -.003 ($p = .38$). * $p < .05$.

Picture condition failed to predict moral elevation, although the relatively wide 95% confidence interval for the mean difference in elevation scores across the two conditions (upper bound = .12, lower bound = -.27) suggests the test may not have had the power required to detect the effect under scrutiny. Although the picture condition to which participants had been assigned did not predict their reported level of moral elevation, there was a significant correlation between moral elevation and participants' own ratings of Alex's dominance; $r = .20$, $p < .01$, suggesting that moral elevation was affected by participants' perception of how dominant they perceived the story protagonist to be.

Did story condition, picture condition, sex, or religiosity predict participants' perceptions of the protagonist's social value?

Participants had been asked to complete a 7-item measure of how attractive they thought Alex would be as a social partner (using a 7-point Likert response scale); participant scores for each item were summed, and divided by 7, to give each participant a 'perceived social attractiveness' score between 1 and 7, with higher values indicating that participants found Alex more attractive as a social partner. Mean social attractiveness ratings are given in Table 13.

A second backwards stepwise linear multiple regression was conducted with story condition, picture condition, sex, and religiosity as predictors, and participants' perceptions of the protagonist's social value as the outcome variable. The coefficients from the regression analysis are shown in Table 14.

Table 13
Mean Ratings of Alex's Social Attractiveness by Picture Condition, Story Condition, Sex, and Religiosity

	Group	<i>n</i>	<i>M(SD)</i>
Picture Condition	High	219	5.93 (.87)
	Low	88	5.88 (.81)
Story Condition	High	161	6.03 (.80)
	Low	146	5.78 (.89)
Sex	Male	134	5.72 (.86)
	Female	172	6.06 (.82)
Religious	Yes	203	5.91 (.85)
	No	103	5.95 (.85)

Note: * $p < .05$

Table 14

Regression Coefficients: Cost Condition, Dominance Condition, Sex, and Religiosity as Predictors of Perceived Social Value

Predictors	B	SE B	β
Step 1			
Constant	5.23	.14	
Cost: Low vs. High	.15	.08	.10
Dominance: Low vs. High	.04	.09	.03
Sex: Female vs. Male	-.37	.08	-.24*
Religiosity: No vs. Yes	.06	.12	.03
Step 2			
Constant	5.26	.12	
Cost Condition: Low vs. High	.14	.08	.10
Religiosity: No vs. Yes	-.37	.08	-.24*
Sex: Female vs. Male	.06	.12	.03
Step 3			
Constant	5.31	.07	
Cost Condition: Low vs. High	.14	.08	.10
Sex: Female vs. Male	-.37	.08	-.24*

Note: R^2 for step 1 = .07 ($p < .001$), ΔR^2 for step 2 = .001 ($p = .65$), ΔR^2 for step 3 = .001 ($p = .60$)

Sex and cost emerged as the only two significant predictors, accounting in combination for 7% of the variance in moral elevation. Follow-up t -tests confirmed that there was a significant sex-difference in ratings of Alex's social attractiveness, with women rating Alex as significantly more attractive than men; $t(305) = 3.57$, $p < .001$, Cohen's $d = .40$. There was also a significant difference in ratings of Alex's social value across the story conditions. Participants reading the high cost story rated Alex as significantly more socially attractive than participants who read the low cost version; $t(305) = .32$, $p = .01$.

Although story condition did not predict moral elevation, there was a significant positive correlation between ratings of Alex's social attractiveness and reported strength of moral elevation, $r = .38$, $p < .001$, suggesting that there may be a relationship between story condition and moral elevation, mediated by perceptions of Alex's social attractiveness, but not one large enough to detect statistically. There

was also a significant positive correlation between participants' own ratings of Alex's dominance and their ratings of his social attractiveness; $r = .25, p < .001$.

Discussion

The results of study 4 offered some support for both prediction 4 and prediction 2, although this support was not unambiguous and may have been weakened by methodological flaws.

Contra prediction 4, the picture conditions to which participants were assigned did not significantly affect participants' ratings of the protagonist's partner value or their reported moral elevation. The relatively weak relationship between the picture condition to which participants were assigned and their ratings of Alex's dominance suggests that the stimuli used in study 4 may not have been fit for purpose; the use of computer generated faces may have lowered the ecological validity of the study, making the predicted effect more difficult to detect. However, participants' own ratings of how dominant Alex was did significantly and positively correlate with both reported moral elevation and ratings of how attractive he was as a social partner, offering some support for the idea that benevolent behaviour is more morally elevating when it is performed by an otherwise socially attractive actor.

Contra prediction 2, there was no significant effect of story condition on reported moral elevation. Participants who read about Alex spending 2.5 hours clearing a 30 yard track reported slightly stronger feelings of moral elevation than participants who read about Alex spending 15 minutes clearing a 5 yard driveway, but the difference was not statistically significant. Participants who read the high cost story version did, however, give significantly higher ratings of Alex's social attractiveness, which itself significantly and positively correlated with reported moral elevation, hinting at a positive relationship between cost incurred by the protagonist and strength of feelings of moral elevation. The stories used to contrast high and low cost benevolence in study 4 may have been less effective as manipulations than those used with more success in study 2. Because the difference in cost incurred by the protagonist in the study 4 stories was less extreme than the differences in the study 2 stories, study 4's experimental manipulation may have been too subtle to produce the same effect.

As in study 1, study 4 found a significant sex difference in moral elevation, adding more weight to the supposition that women are prone to experience stronger

feelings of moral elevation than men. There was no evidence, however, that participant religiosity predicted moral elevation.

There was a significant effect of story condition on perceptions of the protagonist's dominance, which was an unanticipated relationship. It seems reasonable to suppose that the additional physical labour carried out by the protagonist in the high cost story version played into participants' perceptions of how physically formidable he was.

Conclusion And General Discussion

Studies 3A, 3B, and 4 aimed to test predictions about recipient- and actor-characteristics which might be expected to moderate the power of a benevolent action to elicit moral elevation.

It was reasoned in Chapter 2 that we gain little pertinent information about another individual's potential value as a social partner when we see her delivering benefits to recipients whose welfare she has an obvious investment in. When there is no clear incentive for an actor to value the welfare of the recipient, we, as unrelated third-party observers, can make a more accurate assessment of the value the actor may be expected to place on our welfare. On that basis it was predicted that moral elevation would be elicited more strongly by benevolent acts where the recipient was a stranger, than those where the recipient bore some important relationship the actor (be it boss, friend, or family member). Studies 3A and 3B, however, found no support at all for this prediction. On its own, the failure of study 3A to find support for prediction 3 may have been attributable to critical design error, but combined with the lack of support from study 3B the results lend substantial weight to the null-hypothesis.

It was also reasoned in Chapter 2 that, because benevolent actors who are also capable of generating positive externalities are *ceteris paribus* more attractive as social-relationship partners than benevolent actors who do not generate positive externalities, moral elevation should be elicited more easily or more strongly by the actions of individuals who possess other attractive partner qualities. It was therefore predicted that the same benevolent action would elicit greater feelings of moral elevation when it was described as having been performed by a more dominant (physically formidable) actor. The results of study 4 partially conformed with this prediction. Although design flaws may have limited the power of the study to detect

the expected effect, there was some evidence that moral elevation is sensitive to other socially attractive characteristics of a benevolent actor.

The studies reported above also allowed for additional tests of predictions 1 and 2. Although study 1 (reported in the previous chapter) failed to find support for the prediction that moral elevation would be sensitive exclusively to others' benefit conferring behaviours (as opposed to others' moral norm-adhering behaviours), study 3B found strong support for prediction 1. Study 3B used a less ambiguous control story than that used in study 1, and the results of study 3B suggested (in line with those reported by Algoe and Haidt (2009)) that participants classified benefit-conferring behaviours as morally elevating, and virtues of character as inspiring.

In contrast, prediction 2, which was supported by the results of study 2 (reported in the previous chapter), found only slim support from study 4. Participants who read the vignette about the actor who incurred the higher proportionate cost to help the recipient did not report significantly stronger feelings of moral elevation than participants who read the vignette about the actor who incurred a lower proportionate cost. There is reason to think that the manipulation used in study 4 may have been less effective than required to detect the effect that was found in study 2.

Taken in conjunction, then, Chapters 3 and 4 provide some evidence that the situation-detecting component of moral elevation functions in the way we would expect a social-value alert mechanism to operate. Moral elevation does seem to be particularly sensitive to displays of benefit-delivery, it seems to respond to the proportionate costs incurred by benevolent actors, and it appears it may be compounded by actors who are otherwise already socially attractive. It does not, however, appear to be sensitive to the nature of the relationship between the actor and the recipient.

Two trends emerged across the results of studies 1-4; female participants and religiously inclined participants seem to report higher feelings of moral elevation than male or non-religious participants. Although the magnitudes of these differences only reached statistical significance in some of the studies (studies 1 and 4 for sex, and 1 and 3A for religion), the means reported in every study trended in the same direction. The studies reported in Chapter 6 will offer further chances to look for the presence of these relationships, and their potential origin and significance will be considered in Chapter 7.

Finally, it should be noted that the participants used across studies 1 – 4 were recruited from online paid-participant platforms (Crowd Flower and Amazon Mechanical Turk). These crowdsourcing sites have become increasingly popular as sources of relatively inexpensive, plentiful, and quickly collectable data, and many research projects have obtained high quality and reliable data using online collection methods (Buhrmester, Kwang, & Gosling, 2011). It is important to bear in mind, however, that issues such as lack of experimental control (Crump, McDonnell, & Gureckis, 2013) and participant non-naiveté (Chandler, Paolacci, Peer, Mueller, & Ratliff, 2015) may dilute or distort the results produced by studies which rely exclusively on online participants. Study 5, reported in Chapter 6, will offer some variation in participant recruitment methodology to rectify this sole reliance on internet-based samples.

Chapter 5: Moral Elevation's Effects

Abstract

In this chapter, consideration is given to the kinds of computational output moral elevation should be expected to generate if the emotion is an adaptation which functions in the manner suggested by either the reputation-management hypothesis or the relationship-building hypothesis. The theory of recalibrational emotions is described, and the particular type of computational element with which these emotions interact – the internal regulatory variable – is introduced. Predictions concerning which internal regulatory variables moral elevation should interact with are offered. It is proposed that experiences of moral elevation should recalibrate an individual's welfare trade-off ratio, making them more willing to cede welfare to others. It is also proposed that experiences of moral elevation should alter the weight placed by an individual on various motivational priorities; specifically, that an individual should become more concerned with behaving benevolently than with demonstrating strict moral norm-adherence, and that individuals should preferentially seek opportunities to demonstrate a willingness to help others who are less able to repay their actions. It is also predicted, in line with the reputation-management hypothesis and contra the relationship-building hypothesis, that elevated individuals should express stronger motivations to display their benevolence to others than to engage in relationship-building activities with the individual whose behaviour elicited the emotion.

Introduction

Having examined the input side of moral elevation for evidence of special design, this chapter moves on to consider what characteristics the output side of the emotion should be expected to exhibit if it is a functionally specialized part of our naturally developing cognitive architecture.

In studies across the moral elevation literature, participants who report stronger feelings of moral elevation have been shown to make larger donations to charity (Aquino et al., 2011; Thomson & Siegel, 2013; Van de Vyver & Abrams, 2015), make higher offers in a dictator game (Aquino et al., 2011), spend more time helping a researcher with a boring task (Schnall et al., 2010), engage in more volunteering activity (Cox, 2010), demonstrate more organizational citizenship behaviours (Vianello et al., 2010), and take a greater interest in mentoring (Thomson et al., 2014). This combination of results provides strong evidence that moral elevation increases an individual's tendency to engage in benevolent behaviour. As with elevation's elicitors, however, there have been no investigations looking for evidence of nuances in the emotion's output which would be predicted by a theory about its function.

The next section will describe in more detail how emotions are thought to work at a computational level, focusing in particular on the working of a sub-category of emotions that we might expect moral elevation to belong to. The remaining sections will make predictions about observable outputs which would indicate that such computational processes are occurring when people experience moral elevation.

Recalibrational Emotions

As described in Chapter 1, evolutionary psychologists understand emotions as cognitive adaptations which operate as superordinate mechanism-coordinating devices. They mobilize various cognitive subroutines, each emotion creating a pattern of cognitive activation which helped our ancestors to solve some particular, recurrent, and relatively complex adaptive problem (Tooby & Cosmides, 1990b, 2008).

Some emotions initiate immediate, reflexive behavioural sequences (e.g., fear, disgust, or sexual arousal). These are emotions which function to guide us through situations where quickly initiating a fixed pattern of action is necessary to

avert some fitness (or opportunity) cost. Other emotions, though, have less precipitous effects; we experience them as more phenomenologically diffuse, some of them may not have definitive associated facial expressions, and they don't generate the same kind of highly specific action routines as emotions like fear. This second class are known as 'recalibrational emotions', and it is this category to which moral elevation would belong.

Recalibrational emotions alter our motivational priorities by updating the values stored in registers in our minds which help us to adaptively regulate our behaviour. These registers are known as 'internal regulatory variables' (IRVs) - a class of computational element described under a recently proposed computational framework for motivation (Tooby et al., 2008). Our minds use IRVs to store information about variables in the world which can alter the fitness consequences of various types of behaviour.

The fitness consequences of any decisions we make (any course of action we take) are dictated by a number of mutable variables, exogenous and endogenous. The likely fitness-outcome of foraging for a particular kind of food, for example, depends (at least) on the calorific value of the food, the current energy needs/reserves of the individual, the availability of alternative sources of nourishment, the physical exertion required to obtain it, the risk exposure, and the opportunities the individual will miss if she invests time in foraging. Each of these variables is subject to change, as food ripens, personal circumstances change, the seasons move on, etc. In order to make fitness promoting decisions, we need to be able to gather information about the magnitudes of variables like these, and use this information to guide our decision making.

IRVs are proposed to execute this function, acting as our minds' update-able information storage units. They enable us to adaptively match our decision making strategies to our changing environment. Our minds are expected to contain an enormous number of IRVs, to facilitate the highly sophisticated context sensitivity of our behaviour (Cosmides & Tooby, 2013).

WTRs (described in Chapter 2) are an example of a summary IRV. They are computed by summing estimates of a range of variables which dictate the likelihood that our welfare and the welfare of another will overlap. Features of individuals which give us reasons to value their welfare, such as their likely degree of genetic relatedness to us, the chance that they will value our welfare in return, or the

potential they hold as a mating opportunity, are stored as IRVs in our minds (a kinship index, a reciprocity index, a sexual value index etc., for each specific other we encounter). These variables are combined to form the WTR we hold for any individual, which in turn feeds into the decision rule our minds use to make choices about welfare ceding or cost imposing behaviours ($C_{\text{Self}} < B_{\text{Other}} \times \text{WTR}_{\text{Other}} / B_{\text{Self}} > C_{\text{Other}} \times \text{WTR}_{\text{Other}}$). A single IRV can stand as input to a variety of decision rules or summary variables. Kinship indices, for example, feed into the computation of our welfare trade-off ratios, and also into our computation of others' sexual value, thereby affecting the output of decision rules which govern our romantic behaviour (Lieberman et al., 2007).

When features of the environment change, these IRVs need to be updated, or recalibrated. For example, when someone else benefits us unexpectedly, increasing the likelihood that a social relationship with them would be valuable to us, we may experience gratitude, which produces an upward recalibration of: our WTR toward the person to whom we are grateful; the value we place on opportunities to associate with the other; and our motivations to communicate our positive feelings about the other to the other (Lim, 2012). In performing these recalibrations, gratitude helps us build or solidify beneficial social relationships.

Recalibration plays a large part in all emotion programs, but it is expected to be the principal or only role played by emotions like guilt, gratitude, shame, and grief (Tooby & Cosmides, 2008). Rather than producing an immediate and very specific sequence of actions (like the startle reflex when we we're given a fright), recalibrational emotions have a more down-stream effect on our behaviour, by altering the parameters that feed into the decision rules which guide the ways we act. Moral elevation fits the phenomenological profile of a recalibrational emotion because it lacks an associated knee-jerk behavioural response (a feature noted, though not satisfactorily explained, by Fredrickson, 2013).

The sections that follow will consider what recalibrational activities moral elevation may perform. Under the relationship-building hypothesis, the recalibrational activities of moral elevation should be geared towards helping us take advantage of an opportunity to form a social relationship with a high-value potential partner. Under the reputation-management hypothesis, on the other hand, elevation should initiate a pattern of recalibration geared towards helping us safeguard our social reputation in the face of threats from others who display attractive partner

qualities. Of the five predictions given below, confirmation of the first four – predictions 5a, 5b, 6a and 6b - would add equal weight to either hypothesis. Prediction 6c, however, offers a chance to differentiate the two.

Up-Regulating Extrinsic Welfare Trade-Off Ratios

The impetus to benefit another individual can be very different depending on whether or not our actions are being observed, and consequently we are expected to calculate two parallel WTRs; an intrinsic one, which guides unobserved behaviours, and an extrinsic one, which guides our public (or potentially public) behaviours (Tooby & Cosmides, 2008). This extrinsic WTR – or more specifically, one of the parameters that feed into its calculation - would be the most obvious candidate IRV for moral elevation to recalibrate.

Given the importance of benevolent signaling behaviour in human social life, one of the inputs affecting our extrinsic WTR must be some kind of ‘signaling value index’ - a register of the probable marginal returns to be made from public displays of welfare ceding. Up-ratcheting this ‘signal value’ would increase an individual’s willingness to benefit others (in observable situations); whether the signal was intended for all and sundry (à la the reputation-management hypothesis) or for one particular target (à la the relationship-building hypothesis) up-ratcheting this index would be a computationally efficient way to adaptively steer our behaviour.

Prediction 5a: Morally elevated individuals will display higher extrinsic WTRs than will non-elevated individuals

Although there is already plenty of evidence that moral elevation increases welfare-ceding behaviour (summarized above), there has so far been no unambiguous demonstration that elevation is actually recalibrating one of our cognitive registers, rather than just producing some unguided impulse to perform a benevolent act. Even if the interaction between moral elevation and WTRs was considered self-evident, extant research results don’t allow us to estimate the magnitude of the effect. One resource allocation decision in isolation (such as one donation to a charity) only allows us to estimate a lower bound for the actor’s welfare trade-off ratio toward the beneficiary. Giving an actor a series of resource allocation choices with different and opposing fitness outcomes for herself and another person,

however, allows us to identify the upper bound of her WTR (Delton, 2010). Study 5, reported in Chapter 6, will employ a research paradigm capable of giving a more granular picture of moral elevation's behavioural output.

Interaction With Other Attractive Partner-Qualities

Some people rely heavily on their generosity to attract and retain partners, while others have different qualities which allow them to rely less on benevolent signaling to gain access to social relationships. Status, resources, sexual value, formidability, etc. (the qualities discussed in the final section of Chapter 2), raise an individual's partner-value irrespective of their generosity, and so we might expect the presence of these qualities to attenuate the effects of moral elevation on an individual's extrinsic WTR. The 'market value' of otherwise socially-attractive individuals is not as threatened by the socially attractive performances of others (reputation-management hypothesis - Fessler & Haley, 2003), and these individuals have less need to rely on behavioural performances of their own in the course of courting potential social partners (relationship-building hypothesis).

Prediction 5b: The effect of moral elevation on extrinsic WTR will be less pronounced in individuals with relatively-high partner value than in those with relatively low-partner value

Motivational Priorities

Finally, if moral elevation acts in a recalibrational capacity, we should expect to see some effects on the motivational priorities of morally elevated individuals. If one of elevation's functions is to increase our engagement in activities which signal our partner value to others (or a specific other), then we should expect elevated individuals to preferentially seek opportunities to engage in activities with the strongest available signal value. In line with the reasoning laid out in Chapter 2, we might expect that:

Prediction 6a: Morally elevated individuals will express stronger motivations to engage in benevolent behaviour than in non-benevolent but morally laudable behaviour

And

Prediction 6b: Morally elevated individuals will express stronger motivations to benefit others whose welfare they have no incentive to value than those whose welfare they are already invested in

(The failure to find any support for prediction 3 may give us pause in backing prediction 6b, but the data to test prediction 6b were collected at the same time as the data testing prediction 3).

The relationship-building hypothesis and the reputation-management hypothesis make different predictions about the relative value to an elevated individual of signaling opportunities and relationship building opportunities. Although the relationship-building hypothesis allows for some signaling behaviour (in order for the elevated individual to attract the person whose behaviour elicited the emotion), it should predict that motivation towards signaling activities would take a back-seat to motivations to actually build a relationship with the socially attractive other. The reputation-management hypothesis, on the other hand, predicts that elevated individuals would show a stronger preference to demonstrate their benevolent qualities to others than the pursue a relationship with the benevolent actor who generated the emotion. Given the stronger *prima facie* case for the reputation-management hypothesis, this thesis prefers the second prediction:

Prediction 6c: Morally elevated individuals will express stronger motivations to engage in benevolent signaling activities than in relationship-building activities

The following chapter will report the results of two studies intended to test whether moral elevation's outputs conform to the expectations of a functionalist hypothesis. Study 5 will look for evidence of moral elevation recalibrating our welfare trade-off ratios, and study 6 will look for evidence of moral elevation recalibrating our motivational priorities. In addition to further exploring the possibility that moral elevation is a functionally integrated part of our cognitive architecture, it is hoped that the investigation of prediction 6c may allow us to establish a clearer picture of precisely what function moral elevation could have been selected to

perform. If this prediction is met, we would have stronger reason to accept the reputation-management hypothesis than the relationship-building hypothesis.

Chapter 6: Empirical Tests of Predictions 5 and 6

Abstract

This chapter reports the results of two studies. Study 5 aimed to test two predictions: that moral elevation will increase an individual's extrinsic welfare trade-off ratio; and that the effects of moral elevation on an individual's welfare trade-off ratio will be attenuated by the individual's relative partner value. Study 6 aimed to test the prediction that elevated individuals would express a particular pattern of motivational priorities. Study 5 found some support for the first prediction, but no support for the second. Study 6 did not find evidence two of the anticipated motivational differences between more and less elevated individuals, but it did find evidence of the difference anticipated specifically by the reputation-management hypothesis, and not by the relationship-building hypothesis.

Study 5: Does Moral Elevation Increase Extrinsic WTR, And If So, Is This Effect Attenuated By High Partner-Value?

Study 5 aimed to test predictions 5a and 5b:

Prediction 5a: Morally elevated individuals will display higher extrinsic WTRs than will non-elevated individuals

Prediction 5b: Individuals with relatively high partner-value will be less prone to experience moral elevation than individuals with relatively low-partner value

Design

Study 5 employed a between-participant experimental design. Participants were asked to watch a short video (either a morally elevating clip or an amusing one), and then asked to complete a measure designed to show their WTR towards an anonymous stranger. In line with prediction 5a, it was anticipated that participants who watched the elevating video would display higher WTRs than participants who watched the amusing video. Participants were also measured for attractiveness and physical formidability (characteristics which raise the partner value of men and women respectively (Bang Peterson, Sell, Tooby, & Cosmides, 2010; Sell et al., 2008)). In line with prediction 5b, it was anticipated that participants who were relatively attractive or formidable would demonstrate lower WTRs than relatively unattractive or physically weak participants, even in the moral elevation condition.

Method

Participants

106 participants, a mix of undergraduate students and staff, were recruited from University College Isle of Man between 01/05/2018 and 08/02/2019. Ethical approval was granted for the study on 27/03/2018. All participants were residents of the Isle of Man. Participant age ranged between 18 and 75 ($M = 32.82$, $SD = 15.40$), and 63% of participants indicated that they held no religious affiliation.

3 participants failed to complete the measure of the study's dependent variable, and were removed from the dataset. A further 15 participants were removed because the WTRs they expressed during completion of the dependent variable measure were too inconsistent to be functionally interpretable. Full criteria for this removal process are detailed in the procedures section below. The remaining N was 88.

Measures and Materials

The video clip used to elicit moral elevation was a 2.23 minute portion of the video *Kindness Boomerang* published by 'Life Vest Inside', which depicts people on a city street performing kind acts for one another in a 'pay-it-forward' fashion, set to moving music. The video clip used as a control was 2.48 minute clip from a BBC David Attenborough documentary *Planet Earth*, depicting the mating dance of a bird of paradise, set to whimsical music.

The measure used to gauge participants' WTRs was developed by Delton (2010). The measure comprises a series of resource-allocation decision making tasks, in which the participant must choose between allocating a sum of money to themselves, or allocating a different sum of money to an other. Figure 7 gives an example.

Throughout the decisions, the ratio of the values to be allocated varies (in the examples given in Figure 3, the Self amount : Other amount ratios are 27:20, 7:20, and 3:20 respectively). By looking for a participant's 'switch point' - the point at which a participant switches from allocating to themselves to allocating to the other - we can establish an upper and lower bound for the individual's WTR towards the other.

The total set of decisions presented to participants is comprised of 6 anchored sub-sets (mixed and combined). Each sub-set is formed of 10 decisions, and within each set, the amount that the other can be allocated is fixed, anchoring variable sums which can be allocated to the self. The anchored amounts are 19, 23, 37, 46, 68, and 75. The amount that the self can be allocated is determined by multiplying the anchors by different (evenly spaced) ratios: -0.35, -0.15, 0.05, 0.25, 0.45, 0.65, 0.85, 1.05, 1.25, and 1.45 (rounded to nearest integer). Participants were asked to assume that each choice was independent of the others, and that the money could

not be shared between themselves and the other (full task instructions included in Appendix E).

Figure 7. Example of Welfare Trade-off Task Decisions

	Self	Other
Decision 1		
Option 1	54p	0p
Option 2	0p	37p
Decision 2		
Option 1	17p	0p
Option 2	0p	37p
Decision 3		
Option 1	9p	0p
Option 2	0p	37p

Figure 7: Participants must choose between option 1 and option 2 in a series of decisions.

Moral elevation was once again measured using the 8 item moral elevation scale from Vianello et al. (2010). Participants were asked to indicate their agreement with scale items using a 7-point Likert scale ranging from 1 “Strongly disagree” to 7 “Strongly agree”; Cronbach’s α was .81.

To measure participant attractiveness, facial photographs were taken using an iPhone 5, mounted on a *Rhodesy* ‘octopus style’ tripod stand. After initial data collection was complete, the facial photographs were rated by 20 Amazon MTurk workers, recruited from the USA. These workers were presented with each participant photograph sequentially, and asked to rate, on a scale from 1 (‘very unattractive’) to 10 (‘very attractive’), how attractive they thought the person in the photograph was compared to other men/women of a similar age. There was good inter-rater agreement; Cronbach’s α = .98.

To gauge participant formidability, flexed bicep circumference was measured using a standard dress-maker’s tape-measure. Participants were asked to flex

whichever arm they preferred (with sleeves removed where possible), and their biceps was measured at the highest point.

Procedure

Participants were seen by the researcher one-on-one in a private room. Participants were alternately placed into either condition 1 or condition 2 (randomized by the order in which they were recruited).

A facial photograph was taken, with any long hair tied back off the face, glasses removed, and the camera mounted on a tripod to standardize picture angle. Participants' flexed bicep circumference was measured as a gauge of their physical formidability. Although more direct measures of strength exist (e.g. hand-grip or weight-lifting ability), flexed bicep circumference (which has previously been found to account for roughly 50% of the variance in weight-lifting ability (Sell et al., 2010)) was the most cost and time efficient measurement option.

Participants were then shown one of the two short video clips. Following the video, each participant was asked to complete a 'decision making task'. Participants were told that the 'other' in the decision options was another study participant with whom they had been randomly and anonymously paired, and who would not be performing the decision making task. Participants were told they were in the role of 'decision-maker', and their decisions would determine how much money they and the 'other' had the opportunity to win. Making the recipient in the WTR task an anonymous stranger was not only a practically useful aspect of the design – it also ensured that the WTR participants were being asked to display was as close to being comprised only of the value held in participants' 'signal value index' as possible. That is to say, offering participants the chance to benefit an anonymous stranger removed the possibility that any other, non-signal related, reasons for a participant to value the welfare of the recipient would confound the resulting WTRs.

To make the stakes real, participants were told that, at the completion of the study, three participant-pairs would be drawn at random, and the money allocated by the 'decision-maker' would be paid out to both parties of the pair in full. After the decision making task, participants were asked to complete the 8 item questionnaire measure of moral elevation.

Participants' WTRs were calculated following the procedure set out by Delton (2010). For each of the 6 anchored decision sets, participants' WTRs were

calculated by finding the switch point – the point at which the participant switches from giving money to the other to taking money for themselves. If more than one possible switch point exists, the switch point assigned to the set was the one with which the greatest number of decisions conformed; i.e., if 5 of the decisions in the set conformed with one available switch point, and 8 of the decisions conformed with the other, the second switch point would be used to determine the WTR expressed in the set. To gauge the consistency of the assigned switch point, each set was given a ‘consistency maximization score’ (CMS) – a number between 1 and 10, indicating how many choices in the set conformed with the switch point assigned to it.

Each participant’s 6 CMS scores were averaged to give a percentage value representing how uniformly they had expressed the WTRs assigned to them. Because a CMS of 70.84% is expected if participants were answering randomly (Delton, 2010), participants with scores lower than 75% were removed – the remaining N was 88. Once all 6 sets were marked, participants were given a ‘perfect consistency score’ (PCS) – a number between 1 and 6, indicating how many of the 6 anchored sets were perfectly consistent with a single switch point (those with a CMS of 10). Because even one decision that is incongruent with the switch point of a set prevents that set from being perfectly consistent, the PCS offers a more stringent test of participants’ WTR consistency.

The attractiveness ratings gathered for each participant were summed and averaged, to give each participant a mean attractiveness score. Data were manually inputted to and analyzed using SPSS.

Results

Initial data inspection revealed a number of outliers. There was one low-end outlier among the WTR scores, and four low-end outliers among the moral elevation scores. All outliers were replaced with the second lowest non-outlying score for that measure. Following these replacements, WTR scores were normally distributed ($D(86) = .09, p = .06$, but elevation scores were negatively skewed ($D(86) = .11, p = .01$). Reverse Log transformation did not improve the shape of the distribution, and so elevation scores were left untransformed; this should be borne in mind when judging the resilience of the following tests. Both attractiveness scores ($D(79) = .07, p = .20$) and bicep circumference ($D(83) = .09, p = .97$) were normally distributed.

Mean scores for moral elevation, WTR and PCS by condition, sex, and religious persuasion are given in Table 15.

Table 15
Mean Moral Elevation, Welfare Trade-off Ratio, and Perfect Consistency Scores by Condition, Sex, and Religiosity

Group	<i>n</i>	<u>Elevation</u>	<u>WTR</u>	<u>PCS</u>
		<i>M(SD)</i>	<i>M(SD)</i>	<i>M(SD)</i>
Condition				
Elevated	40	4.92(.80)	.82(.42)	5.05(1.24)
Control	39	4.07(.93)	.73(.37)	3.90(2.29)
Sex				
Male	41	4.49(.20)	.78(.40)	4.59(1.92)
Female	36	4.54(.92)	.76(.40)	4.37(1.92)
Religious				
Yes	16	4.51(.98)	.76(.38)	4.57(1.79)
No	63	4.53(.87)	.84(.48)	4.13(2.36)
Total	79	4.51(.96)	.77(.40)	4.48(1.91)

A backwards step-wise linear multiple regression analysis was conducted using moral elevation as the dependent variable and participant age, sex, religious affiliation, and condition as the predictor variables. The purpose of this analysis was to check for any demographic differences in moral elevation, and to ensure that the manipulation had had the anticipated effect. Regression coefficients are shown in Table 16.

As anticipated, condition significantly predicted moral elevation, accounting for 17% of the variance in mean elevation scores. None of the demographic predictors emerged as significant.

A second backwards step-wise linear multiple regression analysis was conducted with condition and moral elevation as the predictors, and WTR as the outcome variable. The resulting coefficients are shown in Table 17.

Table 16

Regression Coefficients: Condition, Sex, Age, and Religiosity as Predictors of Moral Elevation

Predictors	B	SE B	β
Step 1			
Constant	4.20	.24	
Condition: Elevated vs. Control	.79	.18	.44*
Sex: Female vs. Male	.03	.18	.02
Age	-.01	.23	-.06
Religiosity: No vs. Yes	.31	.23	.14
Step 2			
Constant	4.22	.23	
Condition: Elevated vs. Control	.79	.18	.44*
Age	.004	.01	-.06
Religiosity: No vs. Yes	.31	.22	.14
Step 3			
Constant	4.11	.14	
Condition: Elevated vs. Control	.77	.18	.43*
Religiosity: No vs. Yes	.27	.22	.13
Step 4			
Constant	4.19	.13	
Condition: Elevated vs. Control	.74	.18	.41*

Note: R^2 for step 1 = .19 ($p = .002$), ΔR^2 for step 2 <.001 ($p = .86$), ΔR^2 for step 3 = -.003 ($p = .57$), ΔR^2 for step 4 = -.02, ($p = .21$). * $p < .05$

Table 17

Regression Coefficients: Condition and Moral Elevation as Predictors of Welfare Trade-off Ratio

Predictors	B	SE B	β
Step 1			
Constant	.57	.24	
Moral Elevation	.02	.06	.05
Condition: Elevated vs. Control	.09	.10	.10
Step 2			
Constant	.67	.07	
Condition: Elevated vs. Control	.10	.09	.12
Step 3			
Constant	.72	.05	

Note: R^2 for step 1 = .02 ($p = .49$), ΔR^2 for step 2 = -.002 ($p = .68$), ΔR^2 for step 3 = -.02 ($p = .26$)

Contra prediction 5a, neither moral elevation nor the condition in which participants were placed predicted participants' mean WTR. As illustrated in Figure 8, the confidence interval for the mean WTR difference between participants in the two conditions was relatively wide (lower bound = -.08, upper bound = .27), indicating that the test may not have had the required power to detect any effect of condition on WTR.

It was also noticeable that the distribution of welfare trade-off ratios, although it passed tests of normality, was not smoothly distributed, with build-ups of extreme scores at either end of the scale (illustrated in Figure 9). Extreme cases were excluded (i.e., leaving only cases with WTRs > 0 and < 1.14), resulting in a substantially more normal distribution ($D(70) = .09$, $p = .20$) – shown in Figure 10.

Figure 8. Mean Welfare Trade-Off Ratios




Figure 8: Mean welfare trade-off ratios expressed by participants in each of the two conditions. Error bars represent 95% confidence intervals.

Figure 9. Distribution of Welfare Trade-Off Ratios

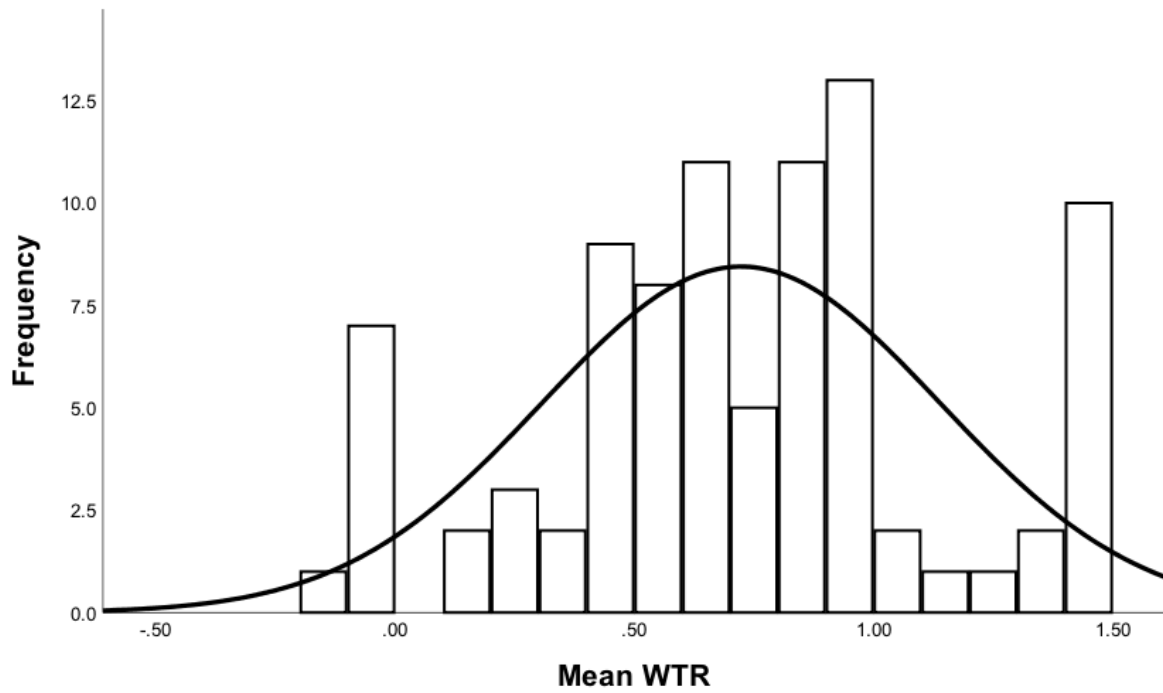


Figure 9: Frequency distribution, with normal curve, of welfare trade-off ratio scores, with build-ups at both extremes of the scale.

Figure 10. Trimmed Distribution of Welfare Trade-Off Ratios




Figure 10: Frequency distribution, with normal curve, of welfare trade-off ratios with extreme cases removed.

With these extreme cases removed, a second backwards step-wise linear multiple regression analysis was conducted with condition and moral elevation as the predictors, and WTR as the outcome variable. The resulting coefficients are shown in Table 18.

Moral elevation did not significantly predict WTR, but there was a significant effect of condition on WTR; $F(1, 66) = 4.67, p = .03, R^2 = .06$. A further simple regression analysis confirmed that, under the trimmed distribution, condition also significantly predicted moral elevation; $F(1, 66) = 18.65, p < .001, R^2 = .22$.

As an exploratory measure, a *t*-test was performed to see whether the difference in perfect consistency scores between participants in the two conditions

was significant. Participants in the elevated condition (under the trimmed distribution) averaged 4.55 perfectly consistent sets out of 6 ($SD = 1.62$), whereas those in the control condition averaged only 3.30 ($SD = 2.47$), a difference which was statistically significant; $t(68) = 2.47, p = .02$. The difference is illustrated in Figure 11.

Table 18

Second Regression Coefficients: Condition and Moral Elevation as Predictors of Welfare Trade-off Ratio

Predictors	B	SE B	β
Step 1			
Constant	.73	.18	
Condition: Elevated vs. Control	.15	.07	.29*
Moral Elevation	-.02	.04	-.07
Step 2			
Constant	.64	.04	
Condition: Elevated vs. Control	.14	.06	.26*

Note: R^2 for step 1 = .07 ($p = .09$), ΔR^2 for step 2 = -.004 ($p = .60$). * $p < .05$.

Figure 11. Mean Perfect Consistency Scores by Condition




Figure 11: Mean number of sets, out of 6, answered perfectly consistently by participants in each of the two conditions.

Both participant attractiveness and bicep circumference were normally distributed; respectively, $D(79) = .07, p = .20$, and $D(83) = .09, p = .97$. A series of simple regression analyses were conducted to explore whether either moral elevation or WTR varied in line with female attractiveness or male bicep circumference. Contra prediction 5b, the results showed no significant relationships: female attractiveness did not predict moral elevation ($R^2 = .01, p = .54$) or WTR ($R^2 = .02, p = .39$); and male bicep circumference did not predict moral elevation ($R^2 = .01, p = .55$) or WTR ($R^2 = .002, p = .76$).

Discussion

Study 5 provided some support for the prediction that more morally elevated individuals would express higher welfare trade-off ratios towards anonymous strangers. Although the regression results with the full set of scores returned no significant results, the sample size for this study was relatively small and the tests may have lacked enough power to detect any present effects. Once the distribution was trimmed of extreme scores, and thereby made more normal, the condition into which participants had been placed did significantly predict the WTR they subsequently displayed. There was also a significant positive correlation between WTR and self-reported moral elevation. There was also a significant effect of video condition on participants' perfect consistency scores – participants who watched the elevating video expressed significantly more consistent WTRs than those who watched the control video. This may indicate a stronger pattern of cognitive activation in the welfare trade-off mechanisms of morally elevated participants, which would be anticipated by the hypothesis that moral elevation acts as a recalibrational emotion.

There was no support, however, for the prediction that more physically formidable males and more attractive females would be less prone to the effects of moral elevation than their less formidable or less attractive counterparts. Neither the attractiveness of female participants, nor the bicep circumference of male participants predicted either the level of moral elevation they reported or the WTR they expressed.

There were a number of methodological problems in study 5 which may have inhibited the ability of the study to detect the predicted effects. Because a single researcher conducted this study in an isolated location with a small population, the overall *N* was relatively small, and there was a lack of environmental control in some aspects of the administration procedures. Not all participants wore clothing which allowed for a clean bicep measure, and the location (and therefore lighting conditions) used to photograph participants varied, contributing to potentially confounding measurement error in both bicep circumference and attractiveness. Any attempts to further explore the effects of physical differences on individuals' proneness to experience moral elevation should consider the use of more sophisticated biometric measuring techniques (e.g., (M E Price, Dunn, Hopkins, & Kang, 2012)).

In addition to testing predictions 5a and 5b, study 5 allowed for a replication of Delton's (2010) demonstration that our minds do in fact house WTRs. Delton illustrated that participants are able to produce remarkably consistent answers on the WTR task (with choices which conform very closely to a single ratio switch point) but unable to achieve anything like the same consistency when performing a mathematically identical task with an externally generated ratio (deciding whether sets of monetary comparisons fell above or below a given exchange rate). An astonishing 45.6% of participants in study 5 had perfect consistency scores of 6 (i.e., all six decision-sets had a single switch point); 17.7% had a PCS of 5, and 13.9% had a PCS of 4. It also bears noting that Delton asked his study participants to imagine a specific 'other' as the recipient in their WTR tasks, the present study used an anonymous stranger as the 'other'. The fact that the same high level of choice consistency emerged here as in Delton's work supports the idea that our welfare trade-off ratio estimator mechanism does include a 'signal value' index, capable of motivating us to confer benefits on others even when the values for all other WTR sub-indices are nil.

There was no significant sex difference in moral elevation in study 5, but as with studies 1 - 4, mean scores for women were slightly higher than mean scores for men. Similarly with religiosity differences, although non-significant, mean elevation for religious participants in study 5 were slightly higher than those for non-religious participants.

Study 6: Do Elevated Individuals Show Particular Motivational Preferences?

Study 6 aimed to test predictions 6a, 6b, and 6c:

Prediction 6a: Morally elevated individuals will express stronger motivations to engage in benevolent behaviour than in non-benevolent but morally laudable behaviour

Prediction 6b: Morally elevated individuals will express stronger motivations to benefit others whose welfare they have no incentive to value than those whose welfare they are already invested in

Prediction 6c: Morally elevated individuals will express stronger motivations to engage in benevolent signaling activities than in relationship-building activities

Design

Study 6 employed an unconventional design, opportunistically harvesting participants from two separate studies already reported above (studies 1 and 3A). The motivation for the design was to maximize data collection (and therefore cost) efficiency. At the close of each of these two studies, participants were asked to complete an ipsative questionnaire intended to gauge their motivational preferences. The intention of study 6 is to explore the results of those motivational preference probes; the purpose of reporting these results here, rather than alongside studies 1 and 3A, is to maintain the expositional integrity of considering moral elevation's inputs separately from its outputs.

In line with prediction 6a, it was anticipated that participants with higher moral elevation scores would express a stronger preference for opportunities to behave benevolently than for opportunities to behave morally; in line with prediction 6b, it was anticipated that participants with higher moral elevation scores would express stronger preferences for opportunities to benefit strangers than loved ones; and in line with prediction 6c, it was anticipated that participants with higher moral elevation scores would express stronger preferences for opportunities to engage in signaling behaviour than opportunities to affiliate with the actor who elicited the emotion.

Method

Participants

Participants from study 1 ($N = 134$, age range 18-68, $M = 34.70$, $SD = 10.02$, 44.8% female, recruited from the USA and India), were combined for analyses with participants from study 3A ($N = 103$, age range 22-69, $M = 36.44$, $SD = 10.64$), 43% female, recruited from the USA), giving a total $N = 227$.

Measures

Participants from both samples had previously been asked to read vignettes describing more or less putatively elevating behaviour and to then rate their experiences of moral elevation. Participants in study 1 had been asked to complete an ipsative version of Vianello et al.'s (2010) 8 item moral elevation measure (given in full in Appendix A). Each of the 8 items had three answer choices, and participants were asked to indicate the option which best described how they felt as they were reading the story. One option for each item was taken from the moral elevation scale, one option was a positively-valenced sensation not associated with elevation, and the third option was 'neither'. Participants in study 3A had previously been asked to read two short vignettes, and then complete an 8 item comparative moral elevation measure (again adapted from Vianello et al. (2010)). For each of the 8 aspects of moral elevation picked up by Vianello et al.'s scale, participants were asked which story, if either, had elicited it most strongly in them (the full measure is given in Appendix B).

Participants' motivations were measured using a 9-item forced-choice questionnaire. The questionnaire was comprised of 9 statement pairs, and participants were asked to choose which item in each pair they agreed with most strongly (full list of items and task instructions included in Appendix F).

The first three questions in the 9 item motivation probe asked participants whether they felt more motivated to engage in benevolent behaviour or moral behaviour, or unmotivated to engage in either kind of behaviour. The first question asked participants whether they more strongly endorsed the statement 'I want to make a positive difference in other peoples' lives' (benevolent option) or the statement 'I think it's very important to fulfill my obligations' (moral option). The second question asked participants to choose between 'Being kind is the most important thing in life' (benevolent option) and 'I always try to do what's right' (moral option). The final question offered 'I want to be a more generous person' (benevolent option) and 'I want to be a more moral person' (moral option).

The second three questions in the 9 item motivation probe asked participants whether they felt more motivated to engage in kindly behaviours towards their loved ones or towards strangers. The first item asked participants to choose between 'I want to help someone less fortunate than myself' (strangers) and 'I want to remind my family that I love them' (loved ones). The second item offered 'Everyone

deserves help when they're at their lowest ebb' (strangers) or 'It's important to be there for the people who are there for you' (loved ones), and the third item offered 'I think we should all work to make society fairer for people who get a raw deal' (strangers) or 'I'm looking forward to the next time I get to hang out with my closest friend' (loved ones).

The final three items aimed to tap into whether morally elevated individuals show a greater concern for signaling their partner value to others (though indiscriminate benevolence) or for building a new social relationship with the protagonist who elicited the elevation (through targeted benevolence). The first item asked participants to choose between 'If I was in [protagonist's] position, I'd like to think I would have done the same thing they did' (indiscriminate) and 'If I met someone like [protagonist], I think we'd be really good friends' (targeted). The second item offered 'I'd like to get more involved in projects that help my local community' (indiscriminate) or 'I'd go out of my way to help someone who acted like [protagonist] did' (targeted), and the third item offered 'Reading the story made me want to be a more caring person' (indiscriminate) or 'As I read the story, I wished I could tell [protagonist] what a good person I thought he was' (targeted).

Procedure

Study 1 participants' moral elevation scores were calculated by giving a score of 1 to each response option from the 8-item elevation scale, and a score of 0 to each control response option and to each 'neither' response option, giving a moral elevation score between 0 and 8, with higher scores indicating stronger feelings of moral elevation.

The use of a comparative format for measuring elevation in Study 3A participants was intended to facilitate exploration of prediction 3 (described in Chapter 4), but the inclusion of the 'neither' choice option also allowed the measure to be scored as an indicator of overall moral elevation. As described in Chapter 4, participants were given a score of 1 each time they answered that either story had evoked a component of elevation, and a score of 0 each time they answered that neither story had evoked the sensation in question. This left Study 3A participants with a moral elevation score between 0 and 8, with higher values taken to indicate stronger overall feelings of moral elevation.

Participants' answers to the motivation probes were scored in three groups. The answers to questions 1-3 were scored so that -1 represented the choice of the 'benevolent' option, 1 represented choosing the 'moral' option, and 0 represented choosing 'neither'. Participants' scores for each of the three questions were summed and averaged, to give participants a mean score for benevolent vs. moral motivation (Motivation 1), such that negative values indicate a stronger motivation engage in benevolent behaviour than to demonstrate fealty to moral norms.

For questions 4-6, a score of -1 indicated having chosen the 'loved ones' option, 1 represented the 'strangers' option, and 0 represented 'neither'. The scores were, again, summed and averaged, giving a mean score for helping loved ones vs. helping strangers (Motivation 2), with negative values indicating a stronger preference for the former, and positive values a stronger preference for the latter.

Finally for questions 7-9, -1 indicated the choice of the 'indiscriminate' option, 1 represented the 'targeted option', and 0 represented 'neither'. Mean scores for indiscriminate vs. targeted benevolence (Motivation 3) ran so that negative numbers indicated a stronger motivation for indiscriminate benevolence, and positive numbers indicated a stronger preference for targeted benevolence.

Data were analysed using IMB's SPSS.

Results

A backwards step-wise linear multiple regression, using motivation 1, motivation 2, and motivation 3 as predictors and moral elevation as the outcome variable. The resulting coefficients are shown in Table 19.

Motivation 1 and Motivation 2 failed to predict moral elevation. There was no significant relationship between whether a participant felt more motivated to act benevolently or morally and the level of moral elevation they reported, and there was likewise no relationship between elevation and whether a participant felt more motivated to invest in loved ones or be kind to strangers. Motivation 3, however, did emerge as a significant predictor, accounting for accounting for 22% of the variance in moral elevation scores; for model 3, $F(1, 225) = 32.79$, $p < .001$, $R^2 = .22$. The more morally elevated participants were, the more likely they were to express stronger motivations to engage in indiscriminate benevolence than to engage in benevolence targeted at the elevating actor.

Table 19

Second Regression Coefficients: Motivation Scores as Predictors of Moral Elevation

Predictors	B	SE B	β
Step 1			
Constant	5.33	.25	
Motivation 1: Benevolence vs. Morality	-.26	.25	-.09
Motivation 2: Loved ones vs. Friends	.15	.24	.06
Motivation 3: Indiscriminate vs. Targeted	-.87	.23	-.33*
Step 2			
Constant	5.38	.24	
Motivation 1: Benevolence vs. Morality	-.21	.23	-.07
Motivation 3: Indiscriminate vs. Targeted	.81	.21	.31*
Step 3			
Constant	5.31	.23	
Motivation 3: Indiscriminate vs. Targeted	-.93	.16	.36*

Note: R^2 for step 1 = .13 ($p < .001$), ΔR^2 for step 2 = -.003 ($p = .37$), ΔR^2 for step 3 = -.002 ($p = .52$). * $p \leq .001$.

Discussion

Study 6 failed to find support for predictions 6a and 6b. Morally elevated individuals did not demonstrate stronger preferences for benefit-conferring over moral norm-adhering behaviour, and nor did they express stronger motivation towards helping lower-value social partners than pre-existing reciprocal relationship partners.

Support was found, however, for prediction 6c. Morally elevated individuals expressed stronger motivations to engage in benevolence which would help others more broadly than to engage in behaviours which would help to foster a social relationship with the person whose behaviour elicited the moral elevation. This result lends some weight to the validity of the reputation-management hypothesis over the relationship-building hypothesis.

Omnibus Analysis: Sex and Religiosity Effects Across Studies

In order to explore further the possible sex and religiosity differences in moral elevation, uncovered in the studies reported above, an omnibus analysis was conducted by pooling data from different data-sets. Participants from studies 1, 3A, 3B, 4, and 5 were combined into one data-set; study 2 participants were not included because they had not completed a comparable measure of the dependent variable, moral elevation. Because religiosity was not recorded for participants in study 4, the total N for religiosity tests was 711 (68.4% identified as religious), whereas for sex tests $N = 876$ (52.2% female). Mean standardized moral elevation scores are broken down by sex and religiosity in Table 20.

Table 20

Mean Standardized Moral Elevation Scores by Sex and Religiosity

Group	<i>M</i>	<i>SD</i>
Sex		
Male	-.14	.99
Female	.12	.99
Religious		
Yes	.07	.96
No	-.16	1.07

Participants' moral elevation scores were standardized, to allow for comparison, and a stepwise linear regression analysis was conducted, with sex and religiosity as predictors, and moral elevation as the outcome variable. The resulting coefficients are shown in Table 21.

Both sex and religiosity significantly predicted moral elevation ($F(2,688) = 10.04, p < .001$). The size of the effect was not large, however, with sex accounting for 2% of the variance in moral elevation, and religiosity accounting for 1%.

Table 21

Regression Coefficients: Sex and Religiosity as Predictors of Moral Elevation

Predictors	B	SE B	β
Model 1			
Constant	.34	.05	
Sex: Female vs. Male	-.27	.08	-.14*
Model 2			
Constant	-.02	.08	
Sex: Female vs. Male	-.26	.08	-.13*
Religiosity: No vs. Yes	.21	.08	.10*

Note: R^2 for model 1 = .02 ($p < .001$), R^2 for model 2 = .03 ($p = .01$)

Chapter 7: Conclusions and Future Research Directions

Abstract

The final chapter summarizes the empirical findings presented in Chapters 3, 4, and 5; general conclusions about the implications of these findings are drawn. There is a discussion of the demographic differences in moral elevation found in some of the studies reported above, and the challenges the sex-difference in particular poses for the reputation-management hypothesis are highlighted. Finally, suggestions for future research directions are made; specifically the potential theoretical utility of exploring cross-cultural differences and similarities, and individual differences in moral elevation.

Introduction

The preceding chapters have sought to explore the likelihood that moral elevation is a naturally selected cognitive adaptation. Although interest in the emotion has gathered steam over the past decade, and its potential social utility is clear, there have so far been no attempts to investigate whether the emotion displays functional properties indicative of ‘special design’. Although no functionalist analyses have been undertaken, at least two (albeit somewhat off-hand) suggestions have been made that elevation may be a naturally selected emotion program which plays an important role in our social behaviour regulation.

In what has here been dubbed the relationship-building hypothesis, Haidt and Fredrickson (Fredrickson, 2004; Haidt, 2000) have proposed that moral elevation could function as a fruitful-relationship-building device, responding to cues indicating the presence of a good-quality prospective social relationship partner, and motivating behaviours which would facilitate affiliation with that individual. In what has been referred to throughout as the reputation-management hypothesis, Fessler and Haley (2003) have proposed that moral elevation may function as a reputation-safeguarding-device, responding to cues indicating that the bar for attracting social relationship partners has been raised, and motivating behaviours which would facilitate us meeting that new bar.

These two hypothesis make largely identical predictions about the computational properties moral elevation should be expected to display; broadly, if either account is correct, moral elevation should be expected to respond to indications that others’ have highly attractive social-partner qualities, and it should motivate behaviours which allow us to demonstrate our own social-partner value. Consequently, this thesis focused on tests of whether elevation displays functional properties that would marry with either account. These tests were organized by inputs and outputs.

Evidence That Moral Elevation’s Inputs Conform to the Expectations of a Functionalist Analysis

Chapters 3 and 4 reported the results of studies intended to test the predictions made in Chapter 2 regarding moral elevation’s elicitors. Because others who are willing to confer benefits on others are, other things equal, more attractive as potential social-relationship partners than others who simply have a tendency to

follow norms, it was expected that moral elevation would be elicited exclusively by others' benefit conferring behaviour, and not by behaviours which are morally-laudable (norm-compliant) but non-benevolent (prediction 1). Although study 1 failed to find any support for this prediction, study 3B (which employed a less conflicted control) found good support. Participants who read a story about a young lady displaying exemplary virtues of character did not classify their emotional reaction to the story as morally elevating, but instead characterized the story as 'inspiring'. Participants who read stories about protagonists who deliberately benefitted others, on the other hand, overwhelmingly classified their emotional reaction as morally elevating. This clear-cut differentiation between moral elevation and admiration (the emotion to which being 'inspired' is the motivational concomitant) replicates a result found by Algoe and Haidt (2009); this not only adds weight to the confidence we can have in the construct validity of moral elevation, it also contributes towards addressing the relative dearth of research evidence supporting fine-grained distinctions between positive emotions in general (Strohming et al., 2011).

Because of the importance we attach to others' willingness to trade-off their welfare for others (as opposed to their ability to do so) when we make judgments about their value as potential social-partners (Delton & Robertson, 2012), it was expected that moral elevation would be sensitive to the valuations displayed by others in their benevolent actions (prediction 2). Study 2 found good support for this prediction, with a strong linear trend indicating that the higher the cost : benefit ratio of a benevolent actor's behaviour, the stronger the feelings of moral elevation it provoked. The same strength of support was not present in study 4, but methodological problems which might account for this difference were discussed in Chapter 4.

It was suggested in Chapter 2 that instances in which individuals benefit others when there is a clear incentive for them to do so would carry less pertinent information for third-parties about the actor's likely social-partner value than instances in which individuals benefit others in the absence of any direct incentives. It was predicted that moral elevation would be less strongly elicited in response to stories about protagonists benefitting their friends, family, or bosses than in response to stories about strangers helping strangers. However, neither study 3A nor study 3B found any support for this prediction, suggesting that the scope and sophistication of moral elevation's operations may be less extensive than proposed.

The failure of this prediction highlights the fact that adaptations (and perhaps particularly cognitive adaptations) are seldom perfectly optimally designed (Gardner, 2017; Olive, 1985). Natural selection can only operate on already extant materials, and it operates only through chance mutation, meaning that its products often do not work in exactly the way they would if a Laplacian demon had set out to design them. Being able to imagine an optimally designed mechanism does not mean that natural selection will have had the time or the opportunity to shape it.

The notion that benevolent actors who have other attractive partner qualities should elicit stronger feelings of moral elevation than benevolent actors who do not (prediction 4) did find some support in study 4. Although the manipulation used was not entirely successful, participants did report significantly stronger feelings of moral elevation when they judged the protagonist of the act as more dominant - a quality which would allow the protagonist to confer a greater number of indirect benefits as a social-relationship partner.

Evidence That Moral Elevation's Outputs Conform to the Expectations of a Functionalist Analysis

Chapter 5 made predictions about the observable outputs that moral elevation should be expected to generate if either the relationship-building or reputation-management hypotheses are correct, and Chapter 6 offered empirical tests of these predictions. It was suggested that moral elevation, if it is a naturally selected emotion program, would fall in the camp of 'recalibrational emotions' – emotions which function largely or solely by altering the operating parameters of other cognitive mechanisms (Cosmides & Tooby, 2013; Lim, 2012; Sell et al., 2008). It was predicted, in turn, that if moral elevation does operate as a recalibration device, it would re-organize the configuration of other cognitive mechanisms in such a way as to produce behaviours which would effectively display our attractiveness as prospective social partners to others in our environment.

There is good evidence that our willingness to confer benefits on others is governed by a decision rule which relies on an internal regulatory variable called a welfare trade-off ratio (Delton, 2010), and study 5 found some evidence that moral elevation may act to upwardly recalibrate this IRV. Participants in the morally elevating condition displayed higher WTRs towards an anonymous other than did participants in the control condition. There was also evidence of an increase in

activation in elevated participants' WTR mechanism; more morally elevated participants expressed significantly more consistent WTRs than their less elevated counterparts. These results, in combination with the results from study 2, indicate a level of functional integration between moral elevation and other relevant cognitive devices (specifically our welfare trade-off estimator, and our welfare trade-off calculator) which would seem to decrease the likelihood that elevation is simply a non-functional byproduct of some other cognitive process.

Study 5 also investigated whether the effects of moral elevation on willingness to benefit others may be attenuated by the possession of other qualities which would contribute to an individual's value as a social-partner. No support was found for this prediction, though methodological flaws which may have contributed to this failure were discussed in Chapter 6.

Study 6 looked for evidence of differences between the motivational priorities of elevated and non-elevated participants. Although there was a strong relationship between feelings of moral elevation and expressing motivations to engage in generally socially positive behaviours, there was little support for any nuanced differences between the motivational priorities of relatively more or less elevated participants. There was no evidence of a relationship between moral elevation and prioritizing the performance of benevolent behaviours over moral norm-adhering behaviours, and no evidence of a relationship between elevation and the prioritization of helping strangers over loved ones.

There was evidence, however, for a relationship between elevation and the prioritization of reputation-guarding behaviours over relationship-building behaviours. This result adds to the *prima facie* reasons for preferring the reputation-management hypothesis to the relationship-building hypothesis. The reputation-management hypothesis can more parsimoniously explain the breadth of moral elevation's prosocial motivational output, and the relationship-building account would struggle to account for a prioritization of broadly benevolent behaviour over behaviour which would facilitate building a social relationship with the individual whose behaviour elicited the emotion.

Achieving a better understanding of moral elevation's outputs will require future research to employ a greater number and diversity of behavioural measures. An over-reliance on self-report measures of moral elevation inhibits a proper picture of the full range of the emotion's effects. If the results reported in study 5 turn out to

be replicable, then Delton's (2010) welfare trade-off ratio task would appear to be a promising method for garnering a finely-scaled estimation of the magnitude of elevation's 'prosocial contagion'. Strategically manipulating the types of recipient that morally elevated participants are given opportunities to help may illuminate some of the more specific features of moral elevation's motivational outputs (in a deeper way than study 6 of this thesis was able to achieve).

Evidence of Sex and Religiosity Differences in Moral Elevation

Across the studies reported in the preceding chapters, there was a trend for female participants and participants who expressed some kind of religious affiliation to report stronger feelings of moral elevation than male or non-religious participants. An omnibus analysis revealed statistically significant differences between male and female elevation scores, and between the elevation scores of religious and non-religious participants, although the magnitude of these effects was not large.

The fact that a number of other studies in the moral elevation literature have found the same sex difference (summarized in Pohling & Diessner, 2016) gives reasonable grounds to suspect that there may be some reliable interaction between sex and sensitivity to moral elevation. The reason for this interaction, though, is unclear. A large meta-analysis (Balliet, Li, Macfarlan, & Van Vugt, 2011) found no significant difference between men and women's tendency to behave cooperatively, but it is a firm, and empirically supported prediction of competitive altruism theory that men engage in benevolent signaling behaviour more than women do (Barclay, 2010; McAndrew, 2012; Raihani & Smith, 2015). If moral elevation has been sculpted by the selection pressures induced by a biological marketplace for cooperative relationships, then the expectation should be that men would be more prone to experience it than women. It could be that higher levels of empathy in women (Christov-Moore et al., 2014) or lower levels of emotional expressivity in men (Deng, Chang, Yang, Huo, & Zhou, 2016) contributed to the violation of this expectation.

This apparent sex-difference in moral elevation was uncovered here by chance, rather than as the result of intentional prediction testing. Future research should include some deliberate investigations of male vs. female proneness to elevation, so as to rule out any confounds that may have come from the research methods and materials used in studies to date. It is worth noting, for example, that

all of the benevolent protagonists used in the eliciting scenarios of this thesis were male (an oversight of the researcher, not an intentional design feature). A fruitful area of further research would be to look at whether differences in the eliciting scenarios used to evoke moral elevation can predict differences in male and female experiences of moral elevation. If it did transpire that the sex difference reported here is a stable feature across different eliciting scenarios, then that result would stand as a challenge to the reputation-management hypothesis (though it would be no more easily explained by the relationship-building hypothesis).

A relationship between religious belief and proneness to experiences of moral elevation has also been noted previously in the elevation literature (Landis et al., 2009), but less frequently than the relationship between elevation and sex. Under some interpretations, religious communities maintain their high coherence by demanding signals of commitment to costly norms and rituals (Alcorta & Sosis, 2006), so it could simply be that religious participants were exhibiting their commitment to a religious norm of benevolence (an injunction to help others being a common theme across all major religions). The higher scores of religious participants could have been an artifact of this kind of commitment; some of the items on the moral elevation scales used in the studies above, such as 'I want to do something good for other people', may receive stronger agreement from religious participants regardless of their actual emotional state. The validity of this explanation could be explored in future research with the use of pre- and post-manipulation moral elevation measures.

Religious communities also tend to be more close-knit, socially integrated, and cooperative networks than non-religious communities (Rai & Fiske, 2011; Sanderson, 2008), however, and so another possibility is that individuals who belong to religious communities face a stronger impetus to maintain an appearance as an attractive social-partner. In the face of such increased competition intensity, greater sensitivity to moral elevation - a mechanism that (putatively) regulates the 'signal value' index of our welfare trade-off calculator - would be beneficial (it should be noted that this explanation would only be congruent with the reputation-management hypothesis, and not with the relationship-building hypothesis).

Conclusion

The results presented in this thesis provide some very preliminary support for the proposal that moral elevation is a functionally integrated part of our species typical cognitive architecture. The positive results found in Chapters 3, 4 and 6 suggest that elevation is associated with a systematic pattern of activation and operation, potentially indicative of special design. There were failures in some of the empirical studies reported above (some more attributable to methodological error than others), and so this support should be offered with caution; substantially stronger results are needed before a confident claim can be made that the emotion is a cognitive adaptation. The studies presented above are by no means comprehensive or water-tight tests of the hypothesis that moral elevation is a naturally selected cognitive adaptation, but nonetheless, there were enough positive findings to warrant future research.

Cross-cultural comparisons

Studies in the extant moral elevation literature have relied on a relatively culturally homogeneous participant pool. The results from study 1 of this thesis, which recruited participants from both the USA and India, went a small distance towards addressing this lack of cross-cultural comparison, but future research will need to significantly broaden and deepen the search for ecological variation in the experience and expression of moral elevation. In order to establish whether moral elevation is truly a cognitive adaptation, and if it is exactly what form it takes, we need to know whether it is experienced to the same degree, by the same types of people, in response to the same stimuli, and with the same cognitive and behavioural outcomes across the variety of social and environmental circumstances that human populations exist in. This can only be achieved if efforts are made by researchers to recruit participants from more geographically and socioeconomically diverse populations.

The results from study 1 of this thesis may strengthen our confidence that moral elevation is not a phenomenon isolated to individuals in WEIRD populations (Western, educated, industrialized, rich, and democratic; Henrich, Heine, & Norenzayan, 2010) – participants from India in fact reported significantly stronger feelings of moral elevation than those from the USA – but they do nothing to address the question of whether any specific environmental contingencies may affect the

expression or development of the emotion. For example, it might be expected (if the explanation offered above for the higher expression of moral elevation in religious participants is solidly reasoned) that populations inhabiting ecologies which increase the imperative to cooperate may express stronger and more frequent experiences of moral elevation. Small-scale populations, and populations living in harsher environments, may be more prone to elevation than the estimates derived from economically comfortable and individualistic cultures have so far led us to expect.

Individual differences

The importance of further exploring the apparent sex and religiosity differences in moral elevation was discussed above, but future research should also consider the possibility that other differences between individuals may affect their proneness to experience moral elevation. A variety of variables affect the marginal returns an individual can expect to make from signaling partner value through benevolent behaviour (Clutton-Brock, 2009). If moral elevation has been naturally selected to help regulate this kind of behaviour, we should expect to see its effects heightened or attenuated by the presence or absence of variables that alter an individual's optimal market competition strategy. It was suggested in Chapter 5 of this thesis that individuals who generate positive externalities, and who therefore (by virtue of already being socially-attractive) enjoy lower marginal returns from benevolent signaling, would be less prone to experiences of moral elevation than individuals who are forced to rely more heavily on their prosocial dispositions to attract social partners. Although study 5 failed to find support for this prediction, it may still bear re-testing with more accurate measures and while controlling for the presence of other positive-externality-generating traits beside physical formidability and attractiveness. Intellect, specialized skills, and surplus resources may all make it less worth an individual's while advertising her partner-qualities through benevolent behaviour (Barclay, 2013).

It may also prove fruitful to consider whether any relationship may exist between an individual's proneness to experiences of moral elevation and her particular life-history strategy (Kaplan & Gangestad, 2005). An individual's life-history strategy describes the pattern of bio-energetic resource allocation decisions she makes across the course of her life-time. Strategies are expressed on a continuum which runs from 'fast' (r) to 'slow' (K), where faster strategies are those

better suited to environments in which the shadow of the future is relatively short, and slower strategies are those better suited to environments in which an individual's life-span is likely to be relatively long (Figueredo et al., 2006;(Ellis, Figueredo, Brumbach, & Schlomer, 2009). In human populations, individuals who display relatively high temporal-discounting rates (taken to be a major component of a faster life-history strategy (Figueredo et al., 2006; Griskevicius, Delton, Robertson, & Tybur, 2011) tend to be less cooperative than their slower-discounting counterparts (Curry, Price, & Price, 2008; Harris & Madden, 2002) (although it bears noting that empirical support for this relationship is not unequivocal - Wu et al., 2017). If individuals with faster life-history strategies invest less in cooperation as a fitness-promotion strategy, we may also expect to find the development of moral elevation less pronounced in these individuals.

Finally, there is evidence that frequency-dependent selection may have endowed human populations with polymorphic cooperative types; stable differences between individuals in terms of how cooperatively they behave, which reflect different strategies for fitness-maximization in a complex social environment (Kurzban & Houser, 2005; Murphy, Handgraaf, & Ackerman, 2011). It seems reasonable to expect that moral elevation, if it functions as proposed, may vary across different cooperative types.

Practical and theoretical utility

The social value of continuing research into a cognitive device which increases our willingness to benefit others seems obvious. If moral elevation is an adaptation which causes contagious benevolence, then its operations must follow a predictable pattern, one that it would be possible to precisely discern and usefully exploit. For example, simply understanding that elevation is likely to be recalibrating an individual's extrinsic, and not her intrinsic welfare trade-off ratio, highlights the fact that fully harnessing any effects of intentionally manipulated moral elevation might be optimized by providing people with public opportunities to behave benevolently. Similarly, if moral elevation's benevolent motivational effects are largely mediated by an increase in an individual's welfare trade-off ratio, then efforts to utilize elevation to increase engagement in collective action problems (such as environmental preservation) may be best pursued by framing the consequences of contribution behaviours in terms of the benefits they will produce for identifiable other individuals,

rather than in terms of self-interest or more abstract and distal benefits. Finally, if elevation's effects are more pronounced in smaller or more closely-knit social contexts, where individuals face a stronger pressure to safeguard their reputations, then attempts to harness the positive social effects of the emotion may be best conducted at the most local level practicable.

Further expanding our understanding of this currently relatively obscure emotion would also help us to achieve a clearer view of how our minds have come to be so adept at managing the adaptive challenges posed by the intricate social groups we inhabit. Recent advancements in understanding and explaining the mental mechanisms underpinning our emotions and motivations in computational terms (Cosmides & Tooby, 2013; Delton et al., 2012; Tooby & Cosmides, 2008) are an interesting and richly fertile new direction for evolutionary psychology. Any explorations which contribute to our understanding of the regulatory and recalibrational cognitive processes behind our social behaviours can only help to sharpen our powers to predict nuanced differences in a range of social outcomes.

References

- Adams, M. J., Johnson, S. A., Lefevre, P., Levesque, V., Hayward, V., Andre, T., & Thonnard, J.-L. (2013). Finger pad friction and its role in grip and touch. *Journal of the Royal Society Interface*, *10*(80). doi:doi:10.1098/rsif.2012.0467
- Al-Shawaf, L., Conroy-Beam, D., Asao, K., & Buss, D. (2015). Human emotions: An evolutionary psychological perspective. *Emotion Review*, *8*(2), 173-186. doi:10.1177/1754073914565518
- Alcorta, C. S., & Sosis, R. (2006). Why ritual works: A rejection of the by-product hypothesis. *Behavioral and Brain Sciences*, *29*, 613-614.
- Alexander, R. D. (1987). *The Biology of Moral Systems*. New York: Aldine de Gruyter.
- Algoe, S. B., & Haidt, J. (2009). Witnessing excellence in action: The 'other-praising' emotions of elevation, gratitude, and admiration. *The Journal of Positive Psychology*, *4*(2), 105-127.
- Aquino, K., McFerran, B., & Laven, M. (2011). Moral identity and the experience of moral elevation in response to acts of uncommon goodness. *Personality Processes and Individual Differences*, *100*(4), 703-718.
- Axelrod, R., & Hamilton, W. D. (1981). The evolution of cooperation. *Science*, *221*(4489), 1380-1396.
- Ayala, F. J., & Campbell, C. A. (1974). Frequency-dependent selection. *Annual Review of Ecology and Systematics*, *5*(1), 115-138.
- Balliet, D., Li, N. P., Macfarlan, S. J., & Van Vugt, M. (2011). Sex differences in cooperation: A meta-analytic review of social dilemmas. *Psychological Bulletin*, *137*(6), 881-909.
- Bang Peterson, M., Sell, A., Tooby, J., & Cosmides, L. (2010). Evolutionary psychology and criminal justice: A recalibrational theory of punishment and reconciliation. In H. Høgh-Olesen (Ed.), *Human morality & sociality: evolutionary and comparative perspectives* (pp. 72-132). New York: Pelgrave Macmillan.
- Barclay, P. (2004). Trustworthiness and competitive altruism can also solve the 'tragedy of the commons'. *Evolution and Human Behaviour*, *25*(4), 209-220.
- Barclay, P. (2010). Altruism as a courtship display: Some effects of third-party generosity on audience perceptions. *British Journal of Psychology*, *101*, 123-135.
- Barclay, P. (2013). Strategies for cooperation in biological markets, especially for humans. *Evolution and Human Behaviour*, *34*, 164-175.
- Barclay, P., & Willer, R. (2007). Partner choice creates competitive altruism in humans. *Proceedings of the Royal Society B: Biological Sciences*, *274*, 749-753.
- Barkow, J. H., Cosmides, L., & Tooby, J. (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.
- Barrett, H. C. (2005). Enzymatic computation and cognitive modularity. *Mind & Language*, *20*(3), 259-287.
- Bliege Bird, R., & Power, E. (2015). Prosocial signaling and cooperation among Martu hunters. *Evolution and Human Behaviour*, *36*(5), 387-397.
- Boyd, R., Gintis, H., & Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science*, *328*(5978), 617-620.
- Buhrmester, M., Kwang, T., & Gosling, S. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, *6*(1), 3-5.

- Buss, D., Haselton, M., Shackelford, T., Bleske, A., & Wakefield, J. (1998). Adaptations, exaptations, and spandrels. *American Psychologist*, *53*(5), 533-548.
- Chandler, J., Paolacci, G., Peer, E., Mueller, P., & Ratliff, K. (2015). Using nonnaive participants can reduce effect sizes. *Association for Psychological Science*, *26*(7), 1131-1139.
- Chaudhuri, A. (2011). Sustaining cooperation in laboratory public goods experiments: A selective survey of the literature. *Experimental Economics*, *14*(1), 47-83.
- Christov-Moore, L., Simpson, E. A., Coude, G., Grigaityte, K., Iacoboni, M., & Ferrari, P. F. (2014). Empathy: gender effects in brain and behaviour. *Neuroscience and biobehavioural reviews*, *46*(4), 604-627. doi:10.1016/j.neubiorev.2014.09.001
- Clutton-Brock, T. (2009). Cooperation between non-kin in animal societies. *Nature*, *462*, 51-57. doi:10.1038/nature08366
- Confer, J. C., Easton, J. A., Fleischman, D. S., Goetz, C. D., Lewis, D. M. G., Perilloux, C., & Buss, D. M. (2010). Evolutionary psychology: controversies, questions, prospects, and limitations. *American Psychologist*, *65*(2), 110-126.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, & J. Tooby (Eds.), *The adapted mind: Evolutionary psychology and the generation of culture*. New York: Oxford University Press.
- Cosmides, L., & Tooby, J. (1994). Origins of domain specificity: The evolution of functional organization. In L. Hirschfeld & S. Gelman (Eds.), *Mappying the mind: Domain specificity in cognition and culture*. New York: Cambridge University Press.
- Cosmides, L., & Tooby, J. (2000). Evolutionary psychology and the emotions. In M. Lewis & J. M. Haviland-Jones (Eds.), *Handbook of emotions* (2nd ed.). New York: Guilford.
- Cosmides, L., & Tooby, J. (2013). Evolutionary psychology: New perspectives on cognition and motivation. *Annual Review of Psychology*, *64*, 201-229.
- Cox, K. S. (2010). Elevation predicts domain-specific volunteerism 3 months later. *The Journal of Positive Psychology*, *5*(5), 333-341.
- Crump, M., McDonnell, J., & Gureckis, T. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioural research. *PLoS ONE*, <https://doi.org/10.1371/journal.pone.0057410>.
- Curry, O., Roberts, S., & Dunbar, R. (2013). Altruism in social networks: Evidence for a 'kinship premium'. *British Journal of Psychology*, *104*, 283-295.
- Curry, O. S., Price, M. E., & Price, J. G. (2008). Patience is a virtue: Cooperative people have lower discount rates. *Personality and Individual Differences*, *44*, 778-783.
- Delton, A. (2010). *A psychological calculus for welfare tradeoffs*. (Doctor of Philosophy in Psychology), University of California, Santa Barbara.
- Delton, A., Cosmides, L., Guemo, M., Robertson, T., & Tooby, J. (2012). The psychosemantics of free riding: Dissecting the architecture of a moral concept. *Journal of Personality and Social Psychology*. Retrieved from doi:10.1037/a0027026
- Delton, A., & Robertson, T. (2012). The social cognition of social foraging: Partner selection by underlying valuation. *Evolution and Human Behaviour*, *33*(6), 715-725.
- Deng, Y., Chang, L., Yang, M., Huo, M., & Zhou, R. (2016). Gender differences in emotional response: Inconsistency between experience and expressivity. *PLoS ONE*, *11*(6).
- Ekman, P. (1999). Basic Emotions. In T. Dalgleish & M. J. Powers (Eds.), *Handbook of Cognition and Emotion* (pp. 45-60). Chichester, UK: Wiley and Sons.

- Ekman, P., & Cordaro, D. T. (2011). What is meant by calling emotions basic? *Emotion Review*, 3(4), 364-370. doi:10.1177/1754073911410740
- Ekman, P., & Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology*, 17(2), 124-129.
- Ellis, B. J., Figueredo, A. J., Brumbach, B. H., & Schlomer, G. L. (2009). Fundamental dimensions of environmental risk: The impact of harsh versus unpredictable environments on the evolution and development of life history strategies. *Human Nature*, 20, 204-268. doi:10.1007/s12110-009-9063-7
- Ellithorpe, M. E., Ewoldsen, D. R., & Oliver, M. B. (2014). Elevation (sometimes) increases altruism: Choice and number of outcomes in elevating media effects. *Psychology of Popular Media Culture*, 4(3), 236-250.
- Englander, Z. A., Haidt, J., & Morris, J. P. (2012). Neural basis of moral elevation demonstrated through inter-subject synchronization of cortical activity during free-viewing. *PLoS ONE*, 7(6), e39384. doi:doi:10.1371/journal.pone.0039384
- Erickson, T. M., & Abelson, J. L. (2012). Even the downhearted may be uplifted: Moral elevation in the daily life of clinically depressed and anxious adults. *Journal of Social and Clinical Psychology*, 31(7), 707-728.
- Fehr, E., & Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behaviour*, 25, 63-87.
- Fehr, E., & Gächter, S. (1999). Cooperation and punishment in public goods experiments. *Institute for Empirical Research in Economics Working Paper No. 10, CESinfo Working Paper Series No. 183*. doi:https://ssrn.com/abstract=203194
- Fehrler, A., & Przepiorka, W. (2013). Charitable giving as a signal of trustworthiness: Disentangling the signaling benefits of altruistic acts. *Evolution and Human Behaviour*, 34, 139-145.
- Fessler, D. M. T., & Haley, K. J. (2003). The strategy of affect: Emotions in human cooperation. In P. Hammerstein (Ed.), *Genetic and cultural evolution of cooperation* (pp. 7-36). Cambridge, MA: MIT Press.
- Figueredo, A. J., Vasquez, G., Brumbach, B. H., Schneider, S. M. R., Safcek, J. A., Tal, I. R., . . . Jacobs, W. J. (2006). Consilience and Life History Theory: From genes to brain to reproductive strategy. *Developmental Review*, 26, 243 - 275.
- Fischbacher, U., Gächter, S., & Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters*(71), 397-404.
- Fletcher, J. A., & Doebeli, M. (2009). A simple and general explanation for the evolution of altruism. *Proceedings of the Royal Society B*, 276(1645), 13-19.
- Fredrickson, B. L. (2004). The broaden-and-build theory of positive emotions. *Philosophical Transactions of the Royal Society of London B*, 359, 1367-1377.
- Fredrickson, B. L. (2013). Positive emotions broaden and build. *Advances in Experimental Social Psychology*, 47, 1-53.
- Fredrickson, B. L., & Branigan, C. (2001). Positive emotions. In T. Mayne & G. A. Bonanno (Eds.), *Emotions: Current issues and future directions* (pp. 123-151). New York: Guildford Press.
- Freeman, D., Aquino, K., & McFerran, B. (2009). Overcoming beneficiary race as an impediment to charitable donation: Social dominance orientation, the experience of moral elevation, and donation behaviour. *Personality and Social Psychology Bulletin*, 35(72-84).
- Gächter, S. (2007). Conditional cooperation: Behavioural regularities from the lab and the field and their policy implications. *Psychology and economics: a promising new cross-disciplinary feild (CESinfo Seminar Series)*, 19-50.

- Gardner, A. (2017). The purpose of adaptation. *Interface focus*, 7(5), 2010005. doi:10.1098/rsfs.2017.0005
- Gendron, M., & Feldman Barrett, L. (2009). Reconstructing the past: A century of ideas about emotion in psychology. *Emotion Review*, 1(4), 316-339.
- Gintis, H. (2003). Solving the puzzle of human prosociality. *Rationality and Society*, 15(2), 155-187.
- Griskevicius, V., Delton, A. W., Robertson, T. E., & Tybur, J. M. (2011). Environmental contingency in life history strategies: The influence of mortality and socioeconomic status on reproductive timing. *Journal of Personality and Social Psychology*, 100(2), 241-254.
- Haidt, J. (2000). The positive emotion of elevation. *Prevention & Treatment*, 3(1), No Pagination Specified Article 3c.
- Haidt, J. (2003a). Elevation and the positive psychology of morality. In C. L. M. Keyes & J. Haidt (Eds.), *Flourishing: Positive psychology and the life well-lived* (pp. 275-289). Washington DC: American Psychological Association.
- Haidt, J. (2003b). The moral emotions. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 852-870). Oxford: Oxford University Press.
- Haidt, J. (2010). Wired to be inspired. In D. Keltner, J. Marsh, & J. Smith (Eds.), *The compassionate instinct* (pp. 86-77). New York: W.W. Norton & Company, Inc. .
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65(4), 613-628.
- Hammerstein, P., & Parker, G. A. (1982). The asymmetric war of attrition. *Journal of Theoretical Biology*, 96(4), 647-682.
- Hardy, C., & Van Vugt, M. (2006). Nice guys finish first: The competitive altruism hypothesis. *Personality and Social Psychology Bulletin*, 32(10), 1402-1413.
- Harris, A. C., & Madden, G. J. (2002). Delay discounting and performance on the prisoner's dilemma game. *The Psychological Record*, 52, 429-440.
- Haselton, M. G., & Ketelaar, T. (2006). Irrational emotions or emotional wisdom? In J. P. Forgas (Ed.), *Hearts and minds: Affective influences on social cognition and behavior* (pp. 21-40). New York: Psychology Press.
- Henrich, J., Heine, S., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33, 61-83.
- Herrmann, B., Thoni, C., & Gächter, S. (2008). Antisocial punishment across societies. *Science*, 319, 1362-1367.
- Høgh-Olesen, H. (2010). Homo sapiens - homo socius: A comparative analysis of human mind and kind. In H. Høgh-Olesen (Ed.), *Human morality & sociality: Evolutionary and comparative perspectives*. New York: Pelgrave Macmillan.
- James, W. (1884). What is an emotion? *Mind*, 9(2), 188-295. doi:10.1093/mind/os-IX.34.188
- Jensen, N. H., & Petersen, M. B. (2011). To defer or to stand up? How offender formidability affects third party moral outrage. *Evolutionary Psychology*, 9(1), 118-136.
- Kaplan, H. S., & Gangestad, S. W. (2005). Life history theory and evolutionary psychology. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 68-95). New Jersey: Wiley.
- Katz, L. (Ed.) (2000). *Evolutionary origins of morality: Cross-disciplinary perspectives*. Exeter: Imprint Academic.

- Kelly, D. (2011). *Yuck! The nature and moral significance of disgust*. Cambridge, MA: MIT Press.
- Keltner, D., & Cordaro, D. T. (2015). Understanding multimodal emotional expressions: Recent advances in basic emotion theory. *Emotion Researcher*, 1-17.
- Krebs, D. L. (2008). Morality: An evolutionary account. *Perspectives on Psychological Science*, 3(3), 149-172.
- Krebs, D. L. (2011). *The origins of morality: An evolutionary account*. Oxford: Oxford University Press.
- Kurzban, R., Burton-Chellew, M. N., & West, S. A. (2015). The evolution of altruism in humans. *Annual Review of Psychology*, 66, 575-599.
- Kurzban, R., DeScioli, P., & Fein, D. (2012). Hamilton vs. Kant: Pitting adaptations for altruism against adaptations for moral judgement. *Evolution and Human Behaviour*, 33, 323-333.
- Kurzban, R., & Houser, D. (2005). Experiments investigating cooperative types in humans: A compliment to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences*, 102(5), 1803-1807.
- Lai, C. K., Haidt, J., & Nosek, B. A. (2014). Moral elevation reduces prejudice against gay men. *Cognition And Emotion*, 28(5), 781-794.
- Landis, S. K., Sherman, M. F., Peidmont, R. L., Kirkhart, M. W., Rapp, E. M., & Bike, D. H. (2009). The relation between elevation and self-reported prosocial behaviour: Incremental validity over the Five-Factor model of personality. *The Journal of Positive Psychology*, 4(1), 71-84.
- Lazarus, R. S. (1991). Progress on a cognitive-motivational-relational theory of emotion. *American Psychologist*, 46(8), 819-834. doi:10.1037/0003-066X.46.8.819
- Lieberman, D., Tooby, J., & Cosmides, L. (2007). The architecture of human kin detection. *Nature*, 445, 727-731.
- Lim, J. (2012). *Welfare tradeoff ratios and emotions: Psychological foundations of human reciprocity*. (Doctor of Philosophy in Anthropology), University of California, Santa Barbara.
- Lin-Healy, F., & Small, D. A. (2013). Nice guys finish last and guys in last are nice: the clash between doing well and doing good. *Social Psychological and Personality Science*, 4(692-698).
- Masel, J. (2011). Genetic drift. *Current Biology*, 21(20). doi:doi:10.1016/j.cub.2011.08.007
- McAndrew, F. T. (2012). Is self-sacrificial competitive altruism primarily a male activity? *Evolutionary Psychology*, 10(1), 50-65.
- Monin, B. (2007). Holier than me? Threatening social comparison in the moral domain. *Revue Internationale De Psychologie Sociale*, 20(1), 53 - 68.
- Murphy, R. O., Handgraaf, M. J. J., & Ackerman, K. A. (2011). Measuring social value orientation. *Judgement and Decision Making*, 6(8), 771-781. doi:10.2139/ssrn.1084189
- Nesse, R. M. (2007). Runaway social selection for displays of partner value and altruism. *Biological Theory*, 2(2), 143-155.
- Noë, R., & Hammerstein, P. (1994). Biological markets: Supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behavioral Ecology & Sociobiology*, 35, 1-11.
- Noë, R., & Hammerstein, P. (1995). Biological Markets. *TREE*, 10(8), 336-339.

- Noë, R., Van Schaik, C. P., & Van Hooff, J. (1991). The market effect: An explanation for pay-off asymmetries among collaborating animals. *Ethology and Sociobiology*, *87*, 97-118.
- Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature*, *437*(27), 1291-1298.
- Oatley, K., & Johnson-Laird, P. N. (2014). Cognitive approaches to emotions. *Trends in cognitive sciences*, *18*(3), 134-140.
- Olive, P. J. (1985). Physiological adaptations and the concepts of optimal reproductive strategy and physiological constraint in marine invertebrates. *Symposia of the Society for Experimental Biology*.
- Panchanathan, K., & Boyd, R. (2004). Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature*, *432*, 499-502.
- Piper, W. T., Saslow, L. R., & Saturn, S. R. (2015). Autonomic and prefrontal events during moral elevation. *Biological Psychology*, *108*, 51-55.
- Pohling, R., & Diessner, R. (2016). Moral elevation and moral beauty: A review of the empirical literature. *Review of General Psychology*, *20*(4), 412-425.
- Price, M. E. (2003). Pro-community altruism and social status in a Shuar village. *Human Nature*, *14*(2), 191-208.
- Price, M. E., Dunn, J., Hopkins, S., & Kang, J. (2012). Anthropometric correlates of human anger. *Evolution and Human Behaviour*, *33*(3), 174-181.
- Rai, T. S., & Fiske, A. P. (2011). Moral psychology is relationship regulation: Moral motives for unity, hierarchy, equality, and proportionality. *Psychological Review*, *118*(1), 57-75.
- Raihani, N. J., & Smith, S. (2015). Competitive helping in online giving. *Current Biology*, *25*(9), 1183-1186.
- Roberts, G. (1998). Competitive altruism: From reciprocity to the handicap principle. *Proceedings of the Royal Society London B*, *265*, 427-431.
- Sanderson, S. K. (2008). Adaptation, evolution, and religion. *Religion*, *38*, 141-156.
- Scarantino, A., & de Sousa, R. (2018). Emotion. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2018 ed.): Metaphysics Research Lab, Stanford University. Retrieved from <<https://plato.stanford.edu/archives/win2018/entries/emotion/>>.
- Schachter, S., & Singer, J. E. (1962). Cognitive, social, and physiological determinants of an emotional state. *Psychological Review*, *69*(379-399).
- Schnall, S., & Roper, J. (2012). Elevation puts moral values into action. *Social Psychological and Personality Science*, *3*(373), 373-378.
- Schnall, S., Roper, J., & Fessler, D. M. T. (2010). Elevation leads to altruistic behaviour. *Psychological Science*, *21*(3), 315-320.
- Sell, A. (2005). *Regulating welfare tradeoff ratios: Three tests of an evolutionary-computational model of human anger*. (Doctor of Philosophy in Psychology), University of California, Santa Barbara.
- Sell, A., Bryant, G. A., Cosmides, L., Tooby, J., Sznycer, D., von Rueden, C., . . . Gurven, M. (2010). Adaptations in humans for assessing physical strength from the voice. *Proceedings of the Royal Society B*, *277*, 3509 - 3518.
- Sell, A., Cosmides, L., Tooby, J., Sznycer, D., von Rueden, C., & Gurven, M. (2009). Human adaptations for the visual assessment of strength and fighting ability from the body and face. *Proceedings of the Royal Society B: Biological Sciences*, *276*, 575-584.

- Sell, A., Tooby, J., & Cosmides, L. (2008). Formidability and the logic of human anger. *Proceedings of the National Academy of Sciences, 106*(35), 15073-15078.
- Siegel, J. T., Thomson, A. L., & Navarro, M. A. (2014). Experimentally distinguishing elevation from gratitude: Oh, the morality. *The Journal of Positive Psychology, 9*(5), 414-427.
- Silvers, J., & Haidt, J. (2008). Moral elevation can induce nursing. *Emotion, 8*, 291-295.
- Smith, M. S., Kish, B. J., & Crawford, C. B. (1987). Inheritance of wealth as human kin investment. *Ethology and Sociobiology, 8*, 171-182.
- Strohinger, N., Lewis, R. L., & Meyer, D. E. (2011). Divergent effects of different positive emotions on moral judgement. *Cognition, 119*, 295-300.
- Thomson, A. L., Nakamura, J., Siegel, J. T., & Csikszentmihalyi, M. (2014). Elevation and mentoring: An experimental assessment of causal relations. *The Journal of Positive Psychology, 9*(5), 402-413.
- Thomson, A. L., & Siegel, J. T. (2013). A moral act, elevation, and prosocial behaviour: Moderators of morality. *The Journal of Positive Psychology, 8*(1), 50-64.
- Todorov, A., Dotsch, R., Porter, J., Oosterhof, N., & Falvello, V. (2013). Validation of data-driven computational models of social perception of faces. *Emotion, 13*, 724-738.
- Todorov, A., & Oosterhof, N. N. (2011). Modeling social perception of faces. *Signal Processing Magazine, 28*(117-122).
- Tooby, J., & Cosmides, L. (1990a). On the universality of human nature and the uniqueness of the individual: The role of genetics and adaptation. *Journal of Personality, 58*(1), 17-67.
- Tooby, J., & Cosmides, L. (1990b). The past explains the present: Emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology, 11*, 375-424.
- Tooby, J., & Cosmides, L. (1996). Friendship and the banker's paradox: Other pathways to the evolution of adaptations for altruism. *Proceedings of the British Academy, 88*, 119-143.
- Tooby, J., & Cosmides, L. (2008). The evolutionary psychology of the emotions and their relationship to internal regulatory variables. In M. Lewis, J. M. Haviland-Jones, & L. F. Barrett (Eds.), *Handbook of emotions* (3rd ed., pp. 114-137). New York: Guilford.
- Tooby, J., Cosmides, L., Sell, A., Lieberman, D., & Sznycer, D. (2008). Internal regulatory variables and the design of human motivation: A computational and evolutionary approach. In A. Elliot (Ed.), *Handbook of approach and avoidance motivation*. Mahwah, NJ: Lawrence Erlbaum Associated.
- Trivers, R. (1971). The evolution of reciprocal altruism. *The Quarterly Review of Biology, 46*(1), 35-57.
- Van de Vyver, J., & Abrams, D. (2015). Testing the prosocial effectiveness of the prototypical moral emotions: Elevation increases benevolent behaviours and outrage increases justice behaviours. *Journal of Experimental Social Psychology, 58*, 23-33.
- Van Vugt, M., Roberts, G., & Hardy, C. (2009). Competitive altruism: A theory of reputation-based cooperation in groups. In R. Dunbar & L. Barrett (Eds.), *Oxford handbook of evolutionary psychology* (pp. 531-541). Oxford: Oxford University Press.
- Vianello, M., Galliani, E. M., & Haidt, J. (2010). Elevation at work: The effects of leaders' moral excellence. *The Journal of Positive Psychology, 5*(5), 390-411.

- Wedekind, C., & Braithwaite, V. A. (2002). The long-term benefits of human generosity in indirect reciprocity. *Current Biology*, *12*, 1012-1015.
- West, S. A., Griffin, A. S., & Gardner, A. (2007). Evolutionary explanations for cooperation. *Current Biology*, *17*, 661-672.
- Williams, G. C. (1966). *Adaptation and natural selection*. Princeton, NJ: Princeton University Press.
- Wilson, E. O. (2012). *The social conquest of earth*. New York: Liveright Publishing Corporation.
- Wu, J., Balliet, D., Tybur, J. M., Arai, S., Van Lange, P. A., & Yamagishi, T. (2017). Life history strategy and human cooperation in economic games. *Evolution and Human Behaviour*, *38*(4), 496-505.
- Zahavi, A. (1975). Mate selection - A selection for a handicap. *Journal of Theoretical Biology*, *53*, 205-214.

Appendix A: Study 1 Materials

Study 1 Vignettes

Benevolent vignette

It had been 10 days, and Oshin knew he probably wouldn't make it through another night. They had been over the centre of the mountain range when the propeller had jammed. The plane was light enough that Oshin had been able to glide it away from the sheer rock and aim at a good-sized patch of clear snow, but the angle of the landing was too steep and the impact was hard.

At first, Oshin had clung to the conviction that the two of them would be found. He forced himself to remain certain that help would arrive at any moment – just another hour more and they would be rescued, just another day longer and the search team would arrive. But as the starvation set in, tiredness and fear had drowned out Oshin's hope, and by now he was sure that he would be dead by the time he and Mika were found.

Oshin had known Mika to be a sensible man. He had flown the young researcher out to the remote villages several times before, and each time he'd been impressed by how knowledgeable Mika was. Oshin had shown him before each flight how to brace himself in his seat if the plane ever got into trouble, but in the heat of the moment Mika had panicked and tried to get to the door at the back of the cabin. He had been on his feet when the plane crashed, and the force of the collision injured him badly. When Oshin regained consciousness, he had pulled Mika free and laid him between the seats of the small plane.

On the first day Oshin had been focused on keeping them warm and trying to signal for help. The radio had been destroyed in the crash, but there was a flare gun and two flares in the cockpit. Oshin fired one flare, but there was heavy mist and he feared it wouldn't be seen, so he saved the second one, hoping the sky would clear by morning. Oshin pulled everything they could use to keep warm from the luggage cases and they hunkered down inside the cabin.

On the second morning the mist had lifted, but the second flare was a dud – it hissed from the gun and buried itself in the snow a short distance from the aircraft. By that time, Oshin was light-headed with hunger, so he set about scouring the cabin for food. He emptied every bag, and turned out every pocket of every garment

he found, but his searching yielded very little. He searched the area around the plane, but the frozen ground was empty – nothing moved and nothing grew.

The grim reality of the situation quickly became apparent: Oshin needed food, but Mika, in his weakened state, needed it more. Oshin had always been a deeply kind man, and so to him it didn't even seem like a choice – he knew what he would do. However bad things got, he couldn't let Mika die. He would rather starve.

Oshin passed the hours in and out of sleep, treading the balance between conserving energy and moving around to keep himself warm. He woke Mika every 12 hours or so, to feed him small portions of the snack food he'd found in the luggage. Oshin knew he could last several days without food, so long as he melted enough snow to drink. He found himself chewing the cuff of his heavy cotton overcoat to distract himself, nibbling away at the frayed edge, swallowing a thread at a time.

Ten days had passed slowly. Oshin found it harder to focus, his mind wondering, unsure at times whether he was awake or asleep. The sharp pain of hunger in his stomach had disappeared after the fourth day, replaced by a gnawing emptiness. He ate pieces of paper, chewed scraps of leather torn from clothes, and scraped small patches of lichen from the nearby rocks – anything to keep his hands and his mouth busy. Throughout the ordeal, Oshin's resolve never wavered. He knew that other men had made different decisions in situations like this, and he could understand how the pain and desperation could drive them to it, but he refused to feed himself and let Mika die.

Eventually he lacked the energy to do anything at all, and now he lay curled under the pile of coats and clothes in the cabin. Mika lay next to Oshin on the cabin floor. He was still weak, but the meager rations had been enough to keep him alive. And as Oshin drifted into unconsciousness, he knew he had done everything he could to make sure Mika got home in one piece.

When he woke again, there were hands lifting Oshin's head and legs. A bright light stung his eyes. The men were talking to him as they lifted him onto the stretcher, but the words were muffled by the sound of the helicopter's blades. Though it hurt to talk, Oshin pulled at the sleeve of the paramedic who was guiding his stretcher towards the helicopter – "Mika...", he managed. The paramedic squeezed Oshin's hand reassuringly; "don't worry", he replied, "we've got him".

Non-benevolent moral vignette

It had been 10 days, and Oshin knew he probably wouldn't make it through another night. They had been over the centre of the mountain range when the propeller had jammed. The plane was light enough that Oshin had been able to glide it away from the sheer rock and aim at a good-sized patch of clear snow, but the angle of the landing was too steep and the impact was hard.

At first, Oshin had clung to the conviction that the two of them would be found. He forced himself to remain certain that help would arrive at any moment – just another hour more and they would be rescued, just another day longer and the search team would arrive. But as the starvation set in, tiredness and fear had drowned out Oshin's hope, and by now he was sure that he would be dead by the time he and Mika were found.

Oshin had known Mika to be a sensible man. He had flown the young researcher out to the remote villages several times before, and each time he'd been impressed by how knowledgeable Mika was. Oshin had shown him before each flight how to brace himself in his seat if the plane ever got into trouble, but in the heat of the moment Mika had panicked and tried to get to the door at the back of the cabin. He had been on his feet when the plane crashed, and the force of the collision killed him outright. When Oshin regained consciousness, he had pulled Mika free and laid him under the wing of the small plane.

On the first day Oshin had been focused on keeping himself warm and trying to signal for help. The radio had been destroyed in the crash, but there was a flare gun and two flares in the cockpit. Oshin fired one flare, but there was heavy mist and he feared it wouldn't be seen, so he saved the second one, hoping the sky would clear by morning. Oshin pulled everything he could use to keep warm from the luggage cases and he hunkered down inside the cabin.

On the second morning the mist had lifted, but the second flare was a dud – it hissed from the gun and buried itself in the snow a short distance from the aircraft. By that time, Oshin was light-headed with hunger, so he set about scouring the cabin for food. He emptied every bag, and turned out every pocket of every garment he found, but his searching yielded nothing. He searched the area around the plane, but the frozen ground was empty – nothing moved and nothing grew.

The grim reality of the situation quickly became apparent: Oshin needed food, but literally the only source of calories was Mika's corpse. Oshin had always

been a deeply moral man, and so to him it didn't even seem like a choice – he knew what he would do. However bad things got, he couldn't resort to cannibalism. He would rather starve.

Oshin passed the hours in and out of sleep, treading the balance between conserving energy and moving around to keep himself warm. He left Mika's body lying in the snow under the wing of the plane, his face covered with a shirt and his arms folded across his chest. Oshin knew he could last several days without food, so long as he melted enough snow to drink. He found himself chewing the cuff of his heavy cotton overcoat to distract himself, nibbling away at the frayed edge, swallowing a thread at a time.

Ten days had passed slowly. Oshin found it harder to focus, his mind wondering, unsure at times whether he was awake or asleep. The sharp pain of hunger in his stomach had disappeared after the fourth day, replaced by a gnawing emptiness. He ate pieces of paper, chewed scraps of leather torn from clothes, and scraped small patches of lichen from the nearby rocks – anything to keep his hands and his mouth busy. Throughout the ordeal, Oshin's resolve never wavered. He knew that other men had made different decisions in situations like this, and he could understand how the pain and desperation could drive them to it, but he refused to eat pieces of Mika's body.

Eventually he lacked the energy to do anything at all, and now he lay curled under the pile of coats and clothes in the cabin. Mika lay outside in the snow. He was still as Oshin has left him 10 days ago – arms folded neatly, his face respectfully covered. And as Oshin drifted into unconsciousness, he knew he had done everything he could to make sure Mika got home in one piece.

When he woke again, there were hands lifting Oshin's head and legs. A bright light stung his eyes. The men were talking to him as they lifted him onto the stretcher, but the words were muffled by the sound of the helicopter's blades. Though it hurt to talk, Oshin pulled at the sleeve of the paramedic who was guiding his stretcher towards the helicopter – "Mika...", he managed. The paramedic squeezed Oshin's hand reassuringly; "don't worry", he said, "we've got him".

Moral Elevation Measure and Task Instructions

Please read each of the following pairs of statements. For each pair, please pick the option which best describes how you felt whilst you read the story. If you felt both of the options listed, please pick the option which you felt most strongly. If you felt neither of the sensations listed, please pick the third option, 'neither'.

1) As I read the story, I felt:

- a) More open towards others
- b) Like a strong and capable person
- c) Neither of the above

2) As I read the story, I felt:

- a) Like a better person
- b) Like someone who's loved
- c) Neither of the above

3) As I read the story, I felt:

- a) Warmth in my chest
- b) Tingles down my spine
- c) Neither of the above

4) As I read the story, I felt:

- a) A lump in my throat
- b) A sudden rush of energy
- c) Neither of the above

5) As I read the story, I felt:

- a) My muscles relaxing
- b) The hairs on my neck stand up
- c) Neither of the above

6) As I read the story, I felt:

- a) Like I wanted to do something good for other people
- b) Like I wanted to work on my skills and talents
- c) Neither of the above

7) As I read the story, I felt:

- a) Like I wanted to do the same thing Oshin was doing
- b) Like doing something playful
- c) Neither of the above

8) As I read the story, I felt:

- a) Like I wanted to become a better person
- b) Like I wanted to sit somewhere quiet and contemplate things
- c) Neither of the above

Appendix B: Study 3A Materials

Study 3A Vignettes

Story 1 – ‘A Day on The River’

Version ‘A’ – Incentive present

It was the *West & Sons* corporate team-building day, and the employees had been split into pairs for a river rafting activity. Each pair was given a small inflatable dinghy, paddles, life-vests, and zip-lock waterproof bags for storing their valuables. Sam had been paired with Ian, an imposing man in his mid 50’s – he was one of the company’s most senior executives.

The two of them donned their life-vests, put their phones, watches, wedding rings, and wallets into the bags, and carried the dinghy down to the water’s edge. Ian sat at the front of the dinghy, and Sam took the seat at the rear. They pushed away from the bank with their paddles and set off down-river with the other rafts. Their route started in a calm tributary, but before long they joined the river proper, where the water picked up speed.

The dinghy bounced over the little eddies in the swift water, as the two men used their paddles to keep the front end pointing down-river. When they approached a bend in the river, Ian lent over the edge of the dinghy to dig his paddle in hard. As Ian did this, Sam saw his zip-lock bag of valuables fall from his pocket and drop over the side of the raft. Ian didn’t notice it fall, but Sam acted quickly.

The water was too fast for them to stop the dinghy being carried down-stream using their paddles, but the river was only waist-height. Sam dropped his paddle at his feet and hopped over the side of the dinghy, planting his feet firmly and grabbing the inflatable raft with one hand to hold it against the current. He fished for Ian’s zip-lock bag with his free hand – it was easy to see under the cold, clear water. Ian helped Sam back into the dinghy and thanked him profusely. “You’re welcome”, Sam smiled, as they righted the boat’s course and set off after the other rafters. He would be damp and cold for the rest of the day, but it was worth it to help his boss out.

Version 'B' – No incentive present

It was the *West & Sons* corporate team-building day, and the employees had been split into pairs for a river rafting activity. Each pair was given a small inflatable dinghy, paddles, life-vests, and zip-lock waterproof bags for storing their valuables. Sam had been paired with Ian, a timid man in his early 20's – he was one of the company's newest and most junior employees.

The two of them donned their life-vests, put their phones, watches, wedding rings, and wallets into the bags, and carried the dinghy down to the water's edge. Ian sat at the front of the dinghy, and Sam took the seat at the rear. They pushed away from the bank with their paddles and set off down-river with the other rafts. Their route started in a calm tributary, but before long they joined the river proper, where the water picked up speed.

The dinghy bounced over the little eddies in the swift water, as the two men used their paddles to keep the front end pointing down-river. When they approached a bend in the river, Ian lent over the edge of the dinghy to dig his paddle in hard. As Ian did this, Sam saw his zip-lock bag of valuables fall from his pocket and drop over the side of the raft. Ian didn't notice it fall, but Sam acted quickly.

The water was too fast for them to stop the dinghy being carried down-stream using their paddles, but the river was only waist-height. Sam dropped his paddle at his feet and hopped over the side of the dinghy, planting his feet firmly and grabbing the inflatable raft with one hand to hold it against the current. He fished for Ian's zip-lock bag with his free hand – it was easy to see under the cold, clear water. Ian helped Sam back into the dinghy and thanked him profusely. "You're welcome", Sam smiled, as they righted the boat's course and set off after the other rafters. He would be damp and cold for the rest of the day, but it was worth it to help the new guy out.

Story 2 – 'A Lift in The Rain'

Version 'A' – Incentive present

It was late, it was cold, and the rain was lashing down. Adam was driving home, looking forward to an early night after a long day at work. Windscreen-wipers working overtime, he had to notch up the radio volume to hear it over the pummeling

of raindrops on the roof. He was driving down a long stretch of road that ran through the edge of the industrial district when he saw a figure trudging through the sheets of rain. It was dark, but Adam could see that it was his good friend Fred. He recognized the well-cut clothes and confident stride of his old college room-mate.

Adam slowed the car and pulled up next to the curb, lowering the far-side window. "You want a ride?" he called. Fred nodded vigorously and pulled open the passenger door. He hustled into the seat, closed the window, and turned to Adam – "hey thanks for this" he grinned "I was starting to think I'd have to swim the rest of the way". "No problem", Adam replied, "where you headed?". Fred brushed the water off his face, "to the bus station", he said, "just a half-mile down the road. My car blew a gasket and I had to leave it with a mechanic back there", he pointed back over his shoulder, "so it's a ride home on the number 40 for me". "You moved recently, right? Where's home now?", Adam asked. "I did, yeah", Fred replied, "I'm on the East side now". "Well hey, look", said Adam, "I'll just run you over there. I'm heading that direction, and it's a heck of a night to be stuck on a winding bus route". "Are you sure?", Fred looked grateful and relieved. "Of course", Adam smiled, "no problem".

When he'd dropped Fred off Adam was 15 miles out of his way. His chances of an early night were lost, but he was glad he'd been able to help a friend out when he needed it.

Version 'B' – No incentive present

It was late, it was cold, and the rain was lashing down. Adam was driving home, looking forward to an early night after a long day at work. Windscreen-wipers working overtime, he had to notch up the radio volume to hear it over the pummeling of raindrops on the roof. He was driving down a long stretch of road that ran through the edge of the industrial district when he saw a figure trudging through the sheets of rain. It was dark, but Adam could see that it was a slightly unkempt man in his early thirties. He had the scruffy clothes and hunched gait of someone down on his luck.

Adam slowed the car and pulled up next to the curb, lowering the far-side window. "You want a ride?" he called. The man nodded vigorously and pulled open the passenger door. He hustled into the seat, closed the window, and turned to Adam – "hey thanks for this" he grinned "I was starting to think I'd have to swim the

rest of the way". "No problem", Adam replied, "where you headed?". The man brushed the water off his face, "to the bus station", he said, "just a half-mile down the road. My car blew a gasket and I had to leave it with a mechanic back there", he pointed back over his shoulder, "so it's a ride home on the number 40 for me". "Where's home?", Adam asked. "Over on the East side", the man replied. "Well hey, look", said Adam, "I'll just run you over there. I'm heading that direction, and it's a heck of a night to be stuck on a winding bus route". "Are you sure?", the man looked grateful and relieved. "Of course", Adam smiled, "no problem".

When he'd dropped the man off Adam was 15 miles out of his way. His chances of an early night were lost, but he was glad he'd been able to help a guy out when he needed it.

Moral Elevation Measure

1) Did either story make you feel more open towards others? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

2) Did either story make you feel like a better person? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

3) Did either story make you feel warm in the chest? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

4) Did either story make you feel a lump in your throat? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

5) Did either story make you feel your muscles relaxing? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

6) Did either story make you feel like you want to do something good for other people? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

7) Did either story make you feel like wanting to do the same thing the protagonist did? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

8) Did either story make you feel like you want to become a better person? (If both stories made you feel this, please pick the story which made you feel it more strongly)

- a) Yes, story A
- b) Yes, story B
- c) No, neither story

Appendix C: Study 3B Materials

Study 3B Vignettes

Story 1 – ‘Boy Writes Book’

Version 1.1 – High incentive

Dylan Seigel, 7, would do anything for his brother Jonah. The pair have always been inseparable. Jonah faces a rare battle with glycogen storage disease, or GSD. The condition causes Jonah’s blood-sugar to drop to potentially fatal levels, forcing his parents to follow an inflexible, round-the-clock schedule of feeding Jonah through a stomach tube in order to keep his metabolism stable.

GSD does not have a cure, and to Dylan, that is unacceptable. Back in 2013, when he was 6 years old, Dylan wrote a book to help his brother and others struggling with the disease. But it’s not your typical page-turner; every penny from sales of *Chocolate Bar* goes directly to a lab at the University of Florida College of Medicine, which is working on a cure for the disease.

To date, Dylan has raised more than \$750,000 in the fight against GSD. Before the money began pouring in from book sales, funding for research into the disease was falling flat. “It’s now a reality. It’s not just a dream that these children can be cured”, Dr David Weinstein, who runs the University of Florida lab said. “It’s an amazing thing that he didn’t just have the idea, he followed through and he actually did it”. *Chocolate Bar* has given those suffering from GSD renewed hope for a better tomorrow.

Version 1.2 – Moderate incentive

Dylan Seigel, 7, would do anything for his best friend Jonah Pournazarian. The pair have been inseparable since they met in preschool. Jonah faces a rare battle with glycogen storage disease, or GSD. The condition causes Jonah’s blood-sugar to drop to potentially fatal levels, forcing his parents to follow an inflexible, round-the-clock schedule of feeding Jonah through a stomach tube in order to keep his metabolism stable.

GSD does not have a cure, and to Dylan, that is unacceptable. Back in 2013, when he was 6 years old, Dylan wrote a book to help his friend and others struggling

with the disease. But it's not your typical page-turner; every penny from sales of *Chocolate Bar* goes directly to a lab at the University of Florida College of Medicine, which is working on a cure for the disease.

To date, Dylan has raised more than \$750,000 in the fight against GSD. Before the money began pouring in from book sales, funding for research into the disease was falling flat. "It's now a reality. It's not just a dream that these children can be cured", Dr David Weinstein, who runs the University of Florida lab said. "It's an amazing thing that he didn't just have the idea, he followed through and he actually did it". *Chocolate Bar* has given those suffering from GSD renewed hope for a better tomorrow.

Version 1.3 – Low incentive

Dylan Seigel, 7, would do anything for a good cause. When he heard about Jonah Pournazarian at an inter-school fundraising event, Dylan decided something needed to be done. Jonah faces a rare battle with glycogen storage disease, or GSD. The condition causes Jonah's blood-sugar to drop to potentially fatal levels, forcing his parents to follow an inflexible, round-the-clock schedule of feeding Jonah through a stomach tube in order to keep his metabolism stable.

GSD does not have a cure, and to Dylan, that is unacceptable. Back in 2013, when he was 6 years old, Dylan wrote a book to help Jonah and others struggling with the disease. But it's not your typical page-turner; every penny from sales of *Chocolate Bar* goes directly to a lab at the University of Florida College of Medicine, which is working on a cure for the disease.

To date, Dylan has raised more than \$750,000 in the fight against GSD. Before the money began pouring in from book sales, funding for research into the disease was falling flat. "It's now a reality. It's not just a dream that these children can be cured", Dr David Weinstein, who runs the University of Florida lab said. "It's an amazing thing that he didn't just have the idea, he followed through and he actually did it". *Chocolate Bar* has given those suffering from GSD renewed hope for a better tomorrow.

Story 2 – ‘Marine Carries Boy’

Version 2.1 – High incentive

Spectators at a children’s triathlon in Florida were brought to tears by the sight of a Marine carrying a little boy with a broken prosthetic leg across the finish line.

Bone cancer survivor Ben Baltz, 11, was participating in his third triathlon of the summer when he had an accident with his prosthetic leg Sunday during the final portion of the race.

The Marine who carried Be, Pfc. Matt Baltz, is Ben’s big brother, and happened to be part of a group of marines helping out at the event. He said “I’d been watching Ben throughout the race, and he was doing very well. It came to the final leg of the race where it was a run and I was sitting at the halfway point with people that were passing out water. As he approached the halfway point, his prosthetic failed and he fell”.

Ben had already jumped up and was trying to fix his prosthetic when his brother ran over to see if he needed help. “He said no, he’d finish by himself and he continued trying to fix it. But after a couple of seconds he knew there was something wrong with it and he was going to need a hand. He was going to finish the race no matter what, but I told him to jump on and we’d finish it together”, Matt said.

Version 2.2 – Moderate incentive

Spectators at a children’s triathlon in Florida were brought to tears by the sight of a Marine carrying a little boy with a broken prosthetic leg across the finish line.

Bone cancer survivor Ben Baltz, 11, was participating in his third triathlon of the summer when he had an accident with his prosthetic leg Sunday during the final portion of the race.

The Marine who carried Be, Pfc. Matt Morgan, is a friend of the Baltz’s, and happened to be part of a group of marines helping out at the event. He said “I’d been watching Ben throughout the race, and he was doing very well. It came to the final leg of the race where it was a run and I was sitting at the halfway point with people that were passing out water. As he approached the halfway point, his prosthetic failed and he fell”.

Ben had already jumped up and was trying to fix his prosthetic when Pfc. Morgan ran over to see if he needed help. “He said no, he’d finish by himself and he

continued trying to fix it. But after a couple of seconds he knew there was something wrong with it and he was going to need a hand. He was going to finish the race no matter what, but I told him to jump on and we'd finish it together", Matt said.

Version 2.3 – Low incentive

Spectators at a children's triathlon in Florida were brought to tears by the sight of a Marine carrying a little boy with a broken prosthetic leg across the finish line.

Bone cancer survivor Ben Baltz, 11, was participating in his third triathlon of the summer when he had an accident with his prosthetic leg Sunday during the final portion of the race.

The Marine who carried Be, Pfc. Matt Morgan of San Diego, said "I'd been watching Ben throughout the race, and he was doing very well. It came to the final leg of the race where it was a run and I was sitting at the halfway point with people that were passing out water. As he approached the halfway point, his prosthetic failed and he fell".

Ben had already jumped up and was trying to fix his prosthetic when Pfc. Morgan ran over to see if he needed help. "He said no, he'd finish by himself and he continued trying to fix it. But after a couple of seconds he knew there was something wrong with it and he was going to need a hand. He was going to finish the race no matter what, but I told him to jump on and we'd finish it together", Matt said.

Story 3 – 'CrossFit'

Some of us go to great lengths, inventing a host of excuses, to avoid exercise at all costs. Not Krystal Cantu.

Despite losing an arm in a car accident last August, Cantu, and avid member of Ballistic CrossFit, has continued to immerse herself in a demanding fitness regimen.

Since her accident, Cantu has competed in high-intensity competitions, including the Working Wounded Games, and the Gladiator Rock n'Run. "I wasn't going to be that person crying on the couch, saying 'why me?'" Cantu explained.

Not only has she continued her rigorous regime, she told the Huffington Post that she is much stronger since her accident, surpassing many of her pre-amputation records.

Cantu wasted no time in getting back to her training after losing her arm. A month after the accident, she was back in the gym.

She doesn't think there is ever a reason to give up; "Every day I get stronger, and every day I prove to myself that the human body is an amazing thing when you have the right mind set. Never make excuses. Know that there is someone out there that has it worse than you do".

Moral Elevation Measure

Please indicate, using the scale provided, how strongly you felt each of the following as you were reading the story.

1. A warmth in my chest
2. A lump in my throat
3. My muscles relaxing
4. I felt like a better person
5. I felt a sense of openness towards other people
6. I wanted to behave like the person in this story did
7. I wanted to do something good for other people
8. I wanted to become a better person

Appendix D: Study 4 Materials

Perceived Partner Value Measure

Please indicate, using the scale provided, how much you agree with the following statements.

1. I hope that good things happen to people like Alex
2. If I met Alex, I think I would not like him [*reverse coded item*]
3. Alex is altruistic
4. Alex is kind
5. I would like to spend time socially with Alex
6. Alex would make a bad team player [*reverse coded item*]
7. Alex would make a good friend

Appendix E: Study 5 Materials

Welfare Trade-Off Ratio Task

Task instructions

The following task will ask you to make a series of decisions about how to allocate money between yourself and another person. You have been randomly paired with another study participant, and a coin-toss has put you in the role of the decider. Your fellow participant will not get to make any decisions, but he/she will be affected by the decisions you make. Once the study is complete, 3 participant pairs will be chosen at random and the allocations made by the decision-maker will be paid out in full.

Below is an example of the type of decisions you will be asked to make.

	You	Other	You	Other	You	Other
Option 1	25	0	15	0	-13	0
Option 2	0	20	0	20	0	20

If you choose the first option, a sum of money will be given to **you** and no money will be given to the other person. If you choose the second option, a sum of money will be given to the **other person** and no money will be given to you. In this example, the first decision asks you to choose between 25p for you or 20p for the other person. If you choose the first option, then you get 25p and the other person gets 0p. If you choose the second option, you get 0p and the other person gets 20p. Similarly, the second decision asks you to choose between 15p for you or 20p for the other person.

In the third decision, you must choose between -13p for you or 20p for the other person. If you choose the first option, then you would have to pay 13p and the other person gets zero. Choosing a negative value for yourself means you lose money. If you choose the second option then you get 0p and the other person gets 20p.

As an example, consider the decisions made by a hypothetical decision-maker, shown below.

	You	Other	You	Other	You	Other
Option 1	25	0	15	0	-13	0
Option 2	0	20	0	20	0	20

Based on the choices made in the example above, for the first decision the decision-maker would receive 25p and the other person 0p. In the second decision the decision-maker would receive 0p and the other 20p. In the third decision the decision-maker would have to pay 13p and the other would receive 0p.

Please try to make each of your decisions independently of your other decisions. That is, try not to let any decision you make influence any of the other decisions you make. The identity of the participant with whom you are paired will be kept anonymous, so you **cannot** share any money you receive with the other person, and they **cannot** share any money they receive with you.

In the following pages, you will be faced with a series of these decisions. For each decision, please circle your preferred option. There are no 'right' or 'wrong' answers to the questions. Please respond based on what feels appropriate to you. Work through the decisions one column at a time.

Task decisions

	You	Other
Option 1	10	0
Option 2	0	23
Option 1	39	0
Option 2	0	37
Option 1	20	0
Option 2	0	23
Option1	4	0
Option2	0	75
Option 1	17	0
Option 2	0	37
Option 1	9	0
Option 2	0	19
Option 1	-10	0
Option 2	0	68
Option 1	12	0
Option 2	0	19
Option 1	31	0
Option 2	0	37
Option 1	-7	0
Option 2	0	19

	You	Other
Option 1	16	0
Option 2	0	19
Option 1	9	0
Option 2	0	37
Option 1	71	0
Option 2	0	68
Option 1	-24	0
Option 2	0	68
Option 1	30	0
Option 2	0	46
Option 1	19	0
Option 2	0	75
Option 1	99	0
Option 2	0	68
Option 1	6	0
Option 2	0	23
Option 1	20	0
Option 2	0	19
Option 1	48	0
Option2	0	46

	You	Other
Option 1	54	0
Option 2	0	37
Option 1	49	0
Option 2	0	75
Option 1	-8	0
Option 2	0	23
Option 1	-3	0
Option 2	0	19
Option 1	2	0
Option 2	0	37
Option 1	24	0
Option 2	0	19
Option 1	-11	0
Option 2	0	75
Option 1	39	0
Option 2	0	46
Option 1	94	0
Option 2	0	75
Option 1	-7	0
Option 2	0	46

	You	Other
Option 1	67	0
Option 2	0	46
Option 1	-6	0
Option 2	0	37
Option 1	34	0
Option 2	0	75
Option1	-16	0
Option2	0	46
Option 1	1	0
Option 2	0	23
Option 1	24	0
Option 2	0	23
Option 1	21	0
Option 2	0	46
Option 1	44	0
Option 2	0	68
Option 1	64	0
Option 2	0	75
Option 1	24	0
Option 2	0	37

	You	Other
Option 1	5	0
Option 2	0	19
Option 1	15	0
Option 2	0	23
Option 1	-3	0
Option 2	0	23
Option 1	31	0
Option 2	0	68
Option 1	2	0
Option 2	0	46
Option 1	28	0
Option 2	0	19
Option 1	85	0
Option 2	0	68
Option 1	-26	0
Option 2	0	75
Option 1	46	0
Option 2	0	37
Option 1	17	0
Option2	0	68

	You	Other
Option 1	1	0
Option 2	0	19
Option 1	3	0
Option 2	0	68
Option 1	33	0
Option 2	0	23
Option 1	79	0
Option 2	0	75
Option 1	-13	0
Option 2	0	37
Option 1	109	0
Option 2	0	75
Option 1	58	0
Option 2	0	68
Option 1	58	0
Option 2	0	46
Option 1	29	0
Option 2	0	23
Option 1	12	0
Option 2	0	46

Appendix F: Study 6 Materials

Ipsative Motivation Probes

Question set 1 – morality vs. benevolence

Please read each of the following pairs of statements. For each pair, please pick the option which you agree with more strongly. If you agree with both of the statements, please pick the one you feel more passionately about.

- 1a. I want to make a positive difference in other peoples' lives
- 1b. I think it's very important to fulfill my obligations
- 1c. Neither of the above

- 2a. Being kind is the most important thing in life
- 2b. I always try to do what's right
- 2c. Neither of the above

- 3a. I want to be a more generous person
- 3b. I want to be a more moral person
- 3c. Neither of the above

Question set 2 – benefitting strangers vs. loved ones

Please read each of the following pairs of statements. For each pair, please pick the option which you agree with more strongly. If you agree with both of the statements, please pick the one you feel more passionately about.

- 1a. I want to help someone less fortunate than myself
- 1b. I want to remind my family that I love them
- 1c. Neither of the above

- 2a. Everyone deserves help when they're at their lowest ebb

- 2b. It's important to be there for the people who are there for you
- 2c. Neither of the above

- 3a. I think we should all work to make society fairer for people who get a raw deal
- 3b. I'm looking forward to the next time I get to hang out with my closest friend
- 3c. Neither of the above

Question set 3 – building relationship vs. signaling benevolence

Please read each of the following pairs of statements. For each pair, please pick the option which you agree with more strongly. If you agree with both of the statements, please pick the one you feel more passionately about.

1a. If I was in [Sam or Adam's / Oshin's] position, I'd like to think I would have done the same thing they did

1b. If I met someone like [Sam or Adam / Oshin], I think we'd be really good friends

1c. Neither of the above

2a. I'd like to get more involved in projects that help my local community

2b. I'd go out of my way to help someone who acted like [Sam or Adam/ Oshins] did

2c. Neither of the above

3a. Reading the story made me want to be a more caring person

3b. As I read the stories, I wished I could tell [Sam or Adam / Oshin] what [good people / a good person] I thought [they were / he was]

3c. Neither of the above