

Uncertainty about the impact of social decisions increases prosocial behaviour

Andreas Kappes^{1,5*}, Anne-Marie Nussberger^{2,5}, Nadira S. Faber^{2,3}, Guy Kahane³, Julian Savulescu³ and Molly J. Crockett^{4*}

Uncertainty about how our choices will affect others infuses social life. Past research suggests uncertainty has a negative effect on prosocial behaviour^{1–12} by enabling people to adopt self-serving narratives about their actions^{1,13}. We show that uncertainty does not always promote selfishness. We introduce a distinction between two types of uncertainty that have opposite effects on prosocial behaviour. Previous work focused on outcome uncertainty (uncertainty about whether or not a decision will lead to a particular outcome). However, as soon as people's decisions might have negative consequences for others, there is also impact uncertainty (uncertainty about how others' well-being will be impacted by the negative outcome). Consistent with past research^{1–12}, we found decreased prosocial behaviour under outcome uncertainty. In contrast, prosocial behaviour was increased under impact uncertainty in incentivized economic decisions and hypothetical decisions about infectious disease threats. Perceptions of social norms paralleled the behavioural effects. The effect of impact uncertainty on prosocial behaviour did not depend on the individuation of others or the mere mention of harm, and was stronger when impact uncertainty was made more salient. Our findings offer insights into communicating uncertainty, especially in contexts where prosocial behaviour is paramount, such as responding to infectious disease threats.

We constantly face decisions that might have consequences for others, and when our decisions do affect others we can never be certain about how they will react^{14–16}. For instance, when there is uncertainty about whether a self-serving action will lead to a potentially negative outcome for others¹—even when there is just a small chance that it will⁴—people are much more likely to act selfishly than when uncertainty is absent. Similarly, if people are uncertain about whether their behaviour will deplete a common resource, they are more likely to overharvest^{10,11}. Such decreases in prosociality might occur because uncertainty enables people to adopt self-serving narratives that allow them to behave selfishly while maintaining a positive self-image^{17,18}. Consistent with this idea, when decision outcomes are uncertain, people optimistically underestimate the chance that self-serving behaviour will cause negative outcomes for others, making self-serving behaviour appear more appropriate to oneself^{1,13,19,20}. Perceptions of social norms (that is, shared beliefs about what people should do in a given situation) mirror these results: self-interested behaviour when outcomes are uncertain not only appears appropriate to oneself, but also to others²¹.

We propose that past research on uncertainty and prosocial behaviour has overlooked the possibility that there are different

types of uncertainty that may have distinct effects on prosocial behaviour. Previous work has focused on what we will call outcome uncertainty (the psychological state in which a decision maker lacks knowledge about what outcome will follow from what choice)²². In the context of social decision-making, most past studies have induced uncertainty about whether or not a decision will lead to a negative outcome for others (Fig. 1a). For instance, a person might lack knowledge about whether or not they will infect another co-worker if they go to work while sick with the flu, or whether a donation will actually reach the people in need. However, outcome uncertainty is only one type of uncertainty present in social interactions. As soon as people's decisions might have consequences for others, they may also lack knowledge about how others' well-being will be impacted by the outcomes of those decisions (Fig. 1b). For example, a person might lack knowledge about how badly the flu will impact a co-worker's well-being, or how much a donation will actually improve the welfare of another person. This type of uncertainty, which we will call impact uncertainty, is uncertainty related to how much the well-being of others will be affected by a particular outcome.

Outcome and impact uncertainty may arise in relation to the same event (for example, infecting another person with the flu), but they correspond to different aspects of this event. Outcome uncertainty occurs when a decision maker lacks knowledge about whether an event (for example, infecting another person) will occur following a particular choice (for example, going to work while sick), and as such bears on the decision maker's causal responsibility for the outcome. In contrast, impact uncertainty occurs when a decision maker lacks knowledge about how an event (for example, infecting another person) will impact the well-being of another person (for example, how badly they will suffer from the infection), and thus relates to the welfare of another person. These two dimensions may respectively contribute to assessments of responsibility and harm magnitude, which independently influence moral judgements²³. We note that the conceptual distinction between 'outcomes' and 'impact' does not correspond with standard constructs in decision theory²⁴; while outcome uncertainty indexes uncertainty about states of the world, impact uncertainty indexes uncertainty about subjective utilities over those states. To support our proposition that outcome and impact uncertainty relate to different assessments that might arise in relation to the same event, we demonstrated that laypeople can reliably distinguish between them in real-world scenarios (see Supplementary Methods).

In contrast with outcome uncertainty, little is known about how impact uncertainty affects social behaviour. This is surprising

¹Department of Psychology, City, University of London, London, UK. ²Department of Experimental Psychology, University of Oxford, Oxford, UK. ³Oxford Uehiro Centre for Practical Ethics, University of Oxford, Oxford, UK. ⁴Department of Psychology, Yale University, New Haven, CT, USA. ⁵These authors contributed equally: Andreas Kappes, Anne-Marie Nussberger. *e-mail: Andreas.Kappes@city.ac.uk; molly.crockett@yale.edu

because impact uncertainty is omnipresent in social interactions: people often lack knowledge about how others will be impacted by the outcomes of their decisions, in large part because other people's subjective experiences are often inaccessible^{14,15}.

Previous research has investigated how people predict the preferences of others, such as preferences about birthday presents²⁵, possessions²⁶, financial or romantic advice^{27–31}, and forgiveness³². This work demonstrates that people often struggle to accurately predict the impact of outcomes on others, even for very close others on important matters such as end-of-life care for a terminally ill spouse³³. Moreover, people are at least partially aware of the lack of insight they have into how others will be impacted by an outcome, with the resulting uncertainty inducing stress and anxiety while making decisions for others, and thereafter doubt and guilt over the decisions made^{34,35}. However, it remains unclear how the experience of impact uncertainty affects prosocial behaviour.

If impact uncertainty activates self-serving narratives in a similar way to outcome uncertainty, people may similarly exploit impact uncertainty to justify self-serving behaviour. For example, when deciding whether to share money with a stranger whose income level is unknown, people might optimistically assume the stranger is rich and thus would benefit little from generosity, creating a self-serving justification for being stingy. However, a recent study suggests that impact uncertainty may increase rather than decrease prosociality, perhaps by activating a different set of narratives around protecting others' welfare³⁶. Participants in this study chose between different amounts of money in exchange for different numbers of electric shocks delivered to either themselves or an anonymous other person. Strikingly, most people were more averse to harming others than themselves. There was no outcome uncertainty in this experiment, but many participants explained their behaviour by appealing to impact uncertainty (for example, "I knew what I could handle but I was less sure about the other person and didn't want to be cruel"), and behavioural indices of uncertainty predicted prosocial behaviour³⁶. This suggests that impact uncertainty may induce precautionary social preferences, where people prefer to avoid the worst-case scenario.

Thus, impact uncertainty might activate different narratives from outcome uncertainty, and consequently have different effects on prosocial behaviour. While outcome uncertainty introduces optimistic and self-serving narratives that mitigate personal responsibility, impact uncertainty may lead people to think more about protecting the welfare of potentially vulnerable others, and thereby increase prosocial behaviour. To test these hypotheses, we independently manipulated impact uncertainty and outcome uncertainty within modified dictator games (studies 1–3) and infectious disease scenarios (study 6). To test whether it is indeed impact uncertainty that drives the observed effects, we examined whether simply mentioning the negative impact (study 4) or any type of uncertainty about the other person (study 5) produced similar results. Finally, we tested whether people are reluctant to contemplate the harm they might cause to others, and whether this reluctance can be overcome by making impact uncertainty salient (study 7).

To manipulate impact uncertainty, we varied the information participants received about the people potentially affected by their decisions. Specifically, in the studies that involved dictator games, we varied the information participants received about the income level of the receiver they were paired with (Fig. 2a–c). All participants in the experimental conditions learned that some of the receivers were "near the bottom of the income scale" and "very dependent on the money they earn to help supplement their income to pay for food and shelter", while others were "at the top of the income scale" and would use the money "to earn a bit of extra spending money; for example, to use for entertainment." In the impact uncertainty condition, participants were then told "The receiver might rank near the top of the income scale, or they might rank near the bottom, or somewhere in the middle."




Fig. 1 | A decision might cause a potentially negative outcome for another person, and the negative outcome may have a negligible or large impact on their well-being. a, Outcome uncertainty is uncertainty about whether or not a decision will cause a negative outcome for another person. In the depicted example, the decision to go to work when feeling ill might or might not lead to infecting a colleague. **b**, Impact uncertainty is uncertainty about how badly the other person's well-being will be impacted by the negative outcome. In the depicted example, infecting a colleague with a disease might cause them a great deal of suffering, or the infection might have only a mild effect.

Our main goal was to test whether people exploit impact uncertainty just like outcome uncertainty and use it to license selfishness, or whether impact uncertainty promotes prosocial behaviour instead. We included two different conditions in our experiments to test these competing hypotheses—a certain-rich and a certain-poor condition. In the certain-rich condition, we told participants that "The receiver ranks near the top of the income scale." If participants exploit impact uncertainty to justify self-serving behaviour and optimistically assume their receiver is rich, prosocial behaviour in the impact uncertainty condition should match prosocial behaviour in the certain-rich condition. In the certain-poor condition, we told participants that "The receiver ranks near the bottom of the income scale." If participants adopt precautionary preferences under impact uncertainty—as we suggest—behaviour in the impact uncertainty condition should match behaviour in the certain-poor condition.

We also included a control condition in which participants did not receive any information about income in general, or their receiver's income in particular. This control condition mirrors how previous research has implemented dictator games³⁷, allowing us to examine whether the introduction of impact uncertainty increases or decreases prosocial decision-making compared with the standard used in decades of research. Furthermore, in everyday life, impact uncertainty is most often implicitly present but not explicitly mentioned, similar to our control condition. Hence, this control condition allows us also to observe whether impact uncertainty increases prosocial behaviour compared with the conditions encountered in everyday life.

To manipulate outcome uncertainty in the dictator games, we replicated the methods of a previous study investigating how outcome uncertainty affects prosocial behaviour¹. Participants played either a standard binary dictator game (Fig. 2d), where a self-serving choice led deterministically to a worse outcome for the receiver (no outcome uncertainty), or a hidden information game, where a self-serving choice led probabilistically to a worse outcome for the receiver (outcome uncertainty; Fig. 2e), but participants had the chance to reveal the outcome of the self-serving option beforehand at no cost.

We used generalized linear models to test the main effects of outcome uncertainty, impact uncertainty and their interaction on prosocial decisions. Thereafter, we tested the simple effects of outcome uncertainty and impact uncertainty with chi-squared tests to determine how the conditions affected the distribution of

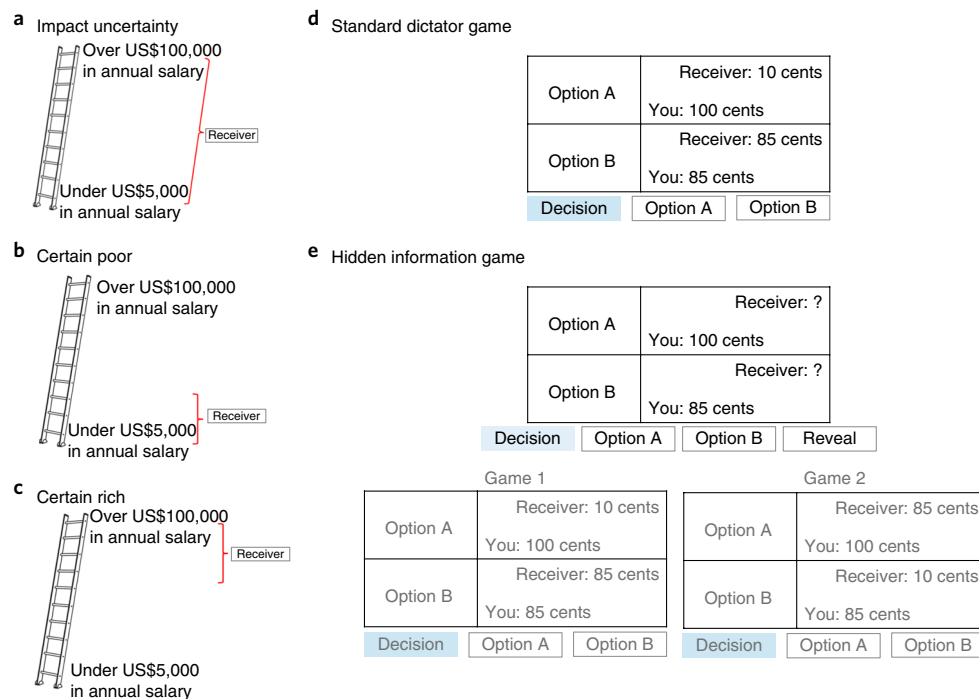


Fig. 2 | Depiction of uncertainty manipulations across experimental conditions. **a**, Participants in the impact uncertainty condition learned that with equal chances, the receiver could be either poor, rich or somewhere in between. **b**, Participants in the certain-poor condition learned that the receiver was poor. **c**, Participants in the certain-rich condition learned that the receiver was rich. **d**, Participants in the outcome certainty condition played a standard binary dictator game. **e**, Participants in the outcome uncertainty condition played a hidden information game. A virtual coin-flip decided whether game 1 or game 2 would be played. In game 1, option A would result in 10 cents for the receiver, but in game 2 it would result in 85 cents. In game 1, option B would result in 85 cents for the receiver, but in game 2 it would result in 10 cents. Participants had the chance to reveal which game was played before making the decision.

prosocial behaviours in our samples. The model predicted decision type (self-serving or prosocial) with separate regressors for outcome uncertainty (type of dictator game: standard or hidden information), impact uncertainty (receiver information: uncertain, certain poor, certain rich or control), and the interaction between outcome and impact uncertainty. As the dependent variable in the standard dictator game, we coded whether participants chose the prosocial option (Fig. 2d, option B). In the hidden dictator game, we coded whether participants chose the reveal option and subsequently took the prosocial option (Fig. 2e, game 1, option B; and Fig. 2e, game 2, option A). We found main effects of outcome uncertainty (χ^2 (d.f. = 1, $n = 832$) = 117.84, $P < 0.001$, Cohen's $d = 0.81$) and impact uncertainty (χ^2 (d.f. = 3, $n = 832$) = 29.33, $P < 0.001$, Cohen's $d = 0.38$), but no interaction (χ^2 (d.f. = 3, $n = 832$) = 2.99, $P = 0.39$). Consistent with previous research^{1-3,5,6,38}, outcome uncertainty reduced prosocial behaviour (Fig. 3a). All of our observed effects remained significant when controlling for participants' income level (see Supplementary Notes).

Next, we examined the effect of impact uncertainty on prosocial behaviour. First, we confirmed that participants were sensitive to the income level of receivers by comparing the proportion of prosocial choices when the receiver had a low income (certain-poor condition) versus high income (certain-rich condition). Indeed, participants in the certain-rich condition were less prosocial than those in the certain-poor condition (χ^2 (d.f. = 3, $n = 417$) = 26.43, $P < 0.001$). To investigate whether this difference was driven by increased generosity towards low-income receivers or decreased generosity towards high-income receivers, we compared each of these conditions with the control condition where participants received no information about the income level of the receivers. As found in previous work³⁹⁻⁴¹, participants in the certain-poor

condition were significantly more prosocial than those in the control condition (χ^2 (d.f. = 1, $n = 418$) = 14.67, $P = 0.001$, Cramer's $V = 0.19$). Meanwhile, participants in the certain-rich condition were not significantly less prosocial than those in the control condition (χ^2 (d.f. = 1, $n = 419$) = 1.81, $P = 0.18$, Cramer's $V = 0.06$). This suggests that the difference in prosocial behaviour between the certain-rich and certain-poor conditions was driven by increased generosity towards low-income receivers.

To test our main prediction that impact uncertainty increases prosocial behaviour, we compared the uncertainty condition with each of the other three conditions. Participants were significantly more prosocial in the uncertainty condition relative to the certain-rich condition (χ^2 (d.f. = 1, $n = 416$) = 14.64, $P = 0.001$, Cramer's $V = 0.19$). These results speak against a self-serving exploitation of impact uncertainty, which would predict that participants assume the receiver in the uncertain condition is rich and thus behave similarly in the uncertain and certain-rich conditions. In contrast, the proportion of prosocial choices in the uncertain condition was not significantly different from that in the certain-poor condition (χ^2 (d.f. = 1, $n = 415$) = 1.80, $P = 0.18$, Cramer's $V = 0.06$). These results suggest that participants in the uncertainty condition erred on the side of caution rather than exploiting uncertainty about the impact on the receiver for their own benefit. Finally, participants in the uncertainty condition were significantly more prosocial than in the control condition (χ^2 (d.f. = 1, $n = 417$) = 6.23, $P = 0.01$, Cramer's $V = 0.12$), suggesting that explicitly mentioning impact uncertainty increases its effect on prosocial behaviour.

One potential alternative explanation for our findings is that participants believed that the average Amazon Mechanical Turk (MTurk) participant was low income. If so, it would be sensible to assume in the uncertainty condition that the receiver is likely to be

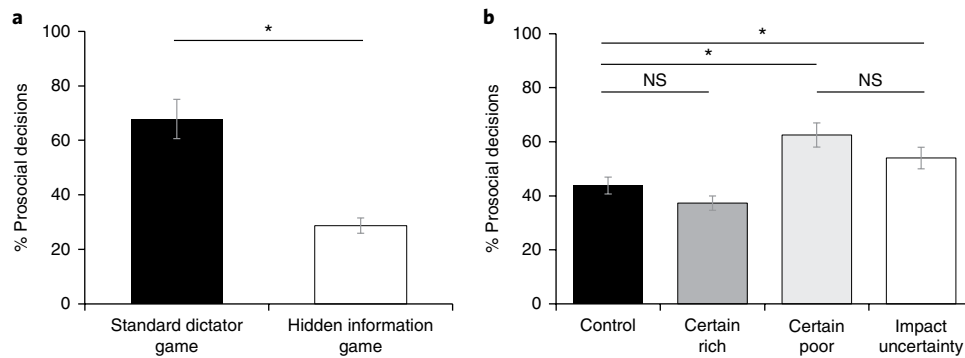


Fig. 3 | Comparison of prosocial decision frequencies for the two game types and four receiver information conditions. a, Percentage of prosocial decisions for the standard dictator game and hidden information game in study 1 ($n=833$). **b**, Percentage of prosocial decisions for the four receiver information conditions (control, certain rich, certain poor and impact uncertainty) in study 1 ($n=833$). Error bars represent s.e. $*P<0.01$. NS, not significant.

low income, and behave accordingly. We ruled out this explanation in two ways. First, we repeated our analysis of impact uncertainty, controlling for participants' beliefs about the average income level of receivers. The effects of impact uncertainty on prosocial behaviour remained significant when controlling for these beliefs, and participants' beliefs about the average income level of receivers did not interact with any of our observed effects (see Supplementary Notes). Second, we conducted a new experiment in which we explicitly controlled participants' beliefs about the income level of the receiver (study 2). In this study, participants in the impact uncertainty condition were instructed: "We pre-selected three receivers: a high-income receiver, a low-income receiver, and a middle-income receiver. You will be paired with one of these receivers at random." Using this belief manipulation, we fully replicated the effects of both outcome and impact uncertainty in study 2 (see Supplementary Notes).

In a third study, we showed that the opposing effects of outcome and impact uncertainty on prosocial behaviour are paralleled at the level of beliefs about social norms, measured via an incentivized coordination game (see Supplementary Notes). Since general beliefs about social norms are independent from specific beliefs about a particular social interaction (including beliefs about the income of the receiver), the fact that participants believe that others think selfish behaviour is wrong when facing impact uncertainty but acceptable when facing outcome uncertainty lends further support to our claim that impact and outcome uncertainty activate distinct narratives about the appropriateness of self-serving behaviour.

However, one might wonder whether impact uncertainty is indeed necessary to increase prosocial behaviour, or whether simply mentioning relative income or the worst possible case is sufficient to produce similar results. In study 4, we tested whether just mentioning the possibility of the receiver being poor, rather than introducing uncertainty about it, is sufficient to increase prosocial behaviour. In this study, the control condition explicitly mentioned that online participants come from all walks of life; some "are poor (meaning their income is below the poverty threshold) and are very dependent on the money they earn on Prolific to help supplement their income to pay for food and shelter". (Prolific is the crowdsourcing platform used to recruit the participants in this study). We induced impact uncertainty by informing participants that the receiver might be either poor or rich, with a probability of 50% each (see Supplementary Methods for details). Again, we found that participants in the impact uncertainty condition behaved more prosocially than participants in the control condition (χ^2 (d.f. = 1, $n=401$) = 5.72, $P=0.018$, Cramer's $V=0.12$) and just as prosocially as participants in the certain-poor condition (χ^2 (d.f. = 1, $n=400$) = 0.035, $P=0.852$, Cramer's $V=0.009$), suggesting that it is

indeed uncertainty about the impact of one's decision that drives the increase in prosocial behaviour.

Study 4 also investigated the question of what level of impact uncertainty may be necessary to enhance prosocial behaviour. In the studies reported so far, participants either faced moderate (33% chance of the receiver being poor) or relatively high (50% chance of the receiver being poor) levels of impact uncertainty. In both cases, participants acted more prosocially relative to the control condition. In study 4, we included a low-impact uncertainty condition, in which participants learned that their receiver might be poor with a 10% chance or rich with a 90% chance. The low-impact uncertainty condition was not significantly different from the control condition ($P=0.284$), leaving open the question of the lowest possible threshold required to elicit impact uncertainty's effect on prosocial behaviour.

Study 4 also included measures of cognitive and affective empathy⁴² and wise reasoning⁴³ to investigate whether impact uncertainty's prosocial effect depends on individual differences. While we found significant main effects of cognitive and affective empathy, as well as wise reasoning, on prosocial behaviour, the conditional effect of impact uncertainty remained significant when we controlled for these individual differences. However, exploratory analyses suggested that the effects of impact uncertainty on prosocial behaviour may be partially mediated by cognitive empathy (see Supplementary Methods for details and further exploratory analyses).

Next, we turned to a potential alternative explanation for our findings: one might argue that it is not necessarily uncertainty about the impact of the outcomes on another person that drives the increase in prosocial behaviour, but rather that uncertainty about any individuating aspect of the receiver might be sufficient. Indeed, previous work has shown that individuation of others increases prosocial behaviour towards them^{44–46}. Thus, reading that the receiver "might be rich, poor, or something in between" in our instructions may have induced participants to think about the receiver as an individual, resulting in increased prosocial behaviour under impact uncertainty. If this was the case, we should observe increased prosocial behaviour under uncertainty even when the dimension participants are uncertain about is irrelevant to the potential harm caused by the outcomes. For instance, it might be sufficient to induce uncertainty about whether or not the receiver is extroverted or introverted.

To test this alternative hypothesis, we provided participants with different information about the receiver's extroversion (study 5). Participants randomly assigned to the uncertain condition read that the receiver "could be extroverted, introverted, or somewhere in between". In the certain-extrovert and certain-introvert conditions, participants read that the receiver is extroverted or introverted,

respectively, while participants in the control condition did not receive any additional information about the receiver. We did not find a significant difference in prosocial decisions across conditions (χ^2 (d.f. = 3, $n = 862$) = 0.94, $P = 0.82$). Thus, participants made as many prosocial choices under impact uncertainty (73%) as they did when being certain about the receiver's intro-/extroversion (both conditions: 72%), or when they did not learn any information about the receiver at all (69%). Taken together, our results show that the increase in prosocial behaviour in our experiments was due to uncertainty about the negative impact of one's actions on others as opposed to simply mentioning negative impact (study 4) or inducing any kind of uncertainty about the other person (study 5).

Next, we examined whether the effects of impact uncertainty are restricted to economic decisions by testing whether we can replicate the results using hypothetical medical decisions concerning the threat of infectious disease. We chose infectious disease since fighting the threat of them depends on behaviours with social consequences (for example, vaccinations, hygiene or isolation). In study 6, participants were asked to imagine the following scenario: "Eight days after you arrived back from a lovely Safari trip to Tanzania, Africa, you feel unwell: feverish and dizzy. You go to the doctor and learn that you have the African flu. The doctor warns you that African flu is contagious: people you come into contact with may get infected. (The following sentence differed across conditions; see below and Supplementary Methods). However, you still feel able to work and you really want to go to the office for finishing a project that is important for your career." (see Supplementary Methods for complete instructions). Participants were then asked to indicate how likely they were to stay home (prosocial intention).

We manipulated impact uncertainty by varying the information participants received about the vulnerability of people they might infect at work. Participants in the impact uncertainty condition read that if they went to work, there was a chance they would infect a young co-worker for whom the African flu would be unproblematic, but also a chance they would infect an older co-worker for whom the African flu would be dangerous. This impact uncertainty condition was compared with a worst-case condition in which participants learned that if they went to work they would infect an older co-worker, as well as with a control condition in which they did not receive additional information about the vulnerability of their co-workers. We found that participants in the impact uncertainty and worst-case conditions were significantly more likely to stay home compared with participants in the control condition ($\chi^2 = 60.75$, $P = 0.004$, $\eta^2 = 0.075$; and $\chi^2 = 90.653$, $P < 0.001$, $\eta^2 = 0.075$, respectively, with χ^2 referring to independent samples Kruskal-Wallis tests and η^2 referring to eta squared). Again, we found no difference in prosocial intention under impact uncertainty compared with the worst-case condition ($P = 0.343$); under impact uncertainty, people formed similar intentions to protect others compared with when the worst case was certain.

Our findings suggest that under impact uncertainty, people consider the potential harmful impact of their actions on others, leading them to err on the side of caution. Yet, people are often reluctant to consider how others could be harmed by their decisions⁴³. It may be that the effects of impact uncertainty rely on overcoming a reluctance to consider harming others, and are only induced if the possibility of others' suffering is made salient to a degree at which people can no longer neglect this possibility when forming a decision. We tested this in a final dictator game study by manipulating the salience of the uncertainty information. We manipulated salience by repeating information, which is one of the most effective ways to increase the salience of information⁴⁷ and the likelihood that people attend to this information^{48,49}.

We included three conditions where the income of the receiver was uncertain, a fact that was made salient to different degrees. In the control condition, participants did not receive any information

about the receiver's income. In the low-salience condition, participants were told that "MTurkers come from all walks of life", additionally mentioning that some "are very dependent on the money they earn on MTurk to help supplement their income" while others "use MTurk as a way to earn a bit of extra spending money". Then, we told them that "We pre-selected three receivers: a high-income receiver, a low-income receiver, and a middle-income receiver. You will be paired with one of these receivers at random." In the high-salience condition, which matched our previous impact uncertainty manipulations, participants received the same information, but we additionally highlighted impact uncertainty by telling participants "The receiver might rank near the top of the income scale, or they might rank near the bottom, or somewhere in the middle." We found a significant difference in the proportion of prosocial decisions across conditions (χ^2 (d.f. = 2, $n = 468$) = 6.02, $P = 0.049$, Cramer's $V = 0.11$). Paired comparisons showed that participants in the high-salience condition were more prosocial than participants in the low-salience condition (χ^2 (d.f. = 1, $n = 314$) = 5.15, $P = 0.023$, Cramer's $V = 0.13$) and control condition (χ^2 (d.f. = 1, $n = 311$) = 4.39, $P = 0.036$, Cramer's $V = 0.12$), which did not differ from one another ($P = 0.86$).

To summarize, we show that uncertainty does not always decrease prosocial behaviour. Instead, the type of uncertainty matters. Replicating previous findings^{1-3,5,6,13,38}, we found that outcome uncertainty (that is, uncertainty about the outcomes of decisions) made people behave more selfishly. However, impact uncertainty about how an outcome would impact another person's well-being increased prosocial behaviour, in economic and health domains. Examining more closely the effect of impact uncertainty on prosocial behaviour, we show that for the increase in prosocial behaviour to occur, simply mentioning negative outcomes or inducing uncertainty about aspects of the other person unrelated to the negative outcome is not sufficient to increase prosocial behaviour. Rather, it seems that uncertainty relating to the impact of negative outcomes on others is needed to increase prosocial behaviour in our studies. Finally, we show that impact uncertainty is only effective when it is salient, thereby potentially overcoming people's reluctance to contemplate the harm they might cause.

Recent theoretical work highlights the power of stories (or narratives) people tell themselves (and others) to justify self-serving behaviour¹⁷. Applied to our findings, this framework suggests that outcome uncertainty activates self-focused narratives that enable people to tell themselves that it is very unlikely that a negative outcome for the other person will occur, allowing them to reap the benefits of self-interested actions without feeling selfish¹⁷. Such self-focused narratives decrease prosocial behaviour by downplaying the potential social costs of self-interested actions. In contrast, our findings suggest that impact uncertainty activates other-focused narratives that include potential social costs, leading participants to adopt behaviours that preserve others' welfare. Notably, such other-focused narratives might also cater for self-image concerns (for example, "only a horrible person would risk infecting a vulnerable other"). Future work, perhaps combining qualitative with quantitative methods, might more directly investigate the content of the narratives motivating people's social behaviour and use these insights to explain how uncertainty encourages (or discourages) prosocial behaviour.

Another important avenue for future research is to examine how other situational features factor into impact uncertainty's effect on prosocial behaviour. We found, for instance, that effect sizes for high and moderate levels of impact uncertainty (50% and 33% chance of negative impact) were similar, whereas the effect size for low-impact uncertainty (10% chance of negative impact) was substantially lower. Based on this observation, we tentatively propose that representations of the expected harmfulness of one's decision's impact on others could be described as a convex function that is

increasingly steep under impact uncertainty: above a certain level, impact uncertainty uniformly affects prosocial behaviour such that people choose whatever option minimizes harm to others. One might speculate that this threshold level depends on further features of the situation. For example, people might be more prone to minimize harm and act prosocially if the harm is physical versus non-physical. People might be willing to maximize their personal outcomes under a 1% chance that another person loses 75 cents, but people might not be willing to maximize their personal outcomes if there is a 1% chance that they could endanger a pregnant woman and the unborn child with an infectious disease.

Our findings highlight the potential for using impact uncertainty to nudge people towards prosocial behaviour. For instance, we found that participants facing impact uncertainty reported they would be more willing to adopt behaviour that would help contain the threat of infectious disease, highlighting the relevance of our findings to addressing global threats. While the communication of such global threats often emphasizes outcome uncertainty (for example, “What are the chances of a devastating pandemic occurring in the next 50 years?”⁵⁰), impact uncertainty is rarely communicated. However, our work suggests that when communicating uncertainty, policy makers, public health officials and others should consider which type of uncertainty they intend to communicate. Since outcome uncertainty biases people towards self-interested behaviour, highlighting impact uncertainty instead may lead to more prosocial decision-making.

Methods

All studies were approved by the University of Oxford's Central University Research Ethics Committee (approval number: MSD-IDREC-C1-2014-005) and participants in each study gave their informed consent beforehand.

Study 1. Participants. We determined sample size using G-Power 3.1 (ref. ⁵¹; see Supplementary Methods). A total of 833 participants were recruited via MTurk. MTurk provides reliable participants who are ethnically and socioeconomically more diverse than university-recruited participants^{52,53}. We paid participants in line with the US minimum wage.

Procedure. Participants were instructed that their decision determined the exact monetary amount they themselves and a receiver would obtain. The experiment did not involve deception as a corresponding number of receivers was randomly recruited from unrelated studies on MTurk and paid according to participants' choices. We manipulated outcome and impact uncertainty as between-subject factors with two and four levels, respectively.

To manipulate outcome uncertainty, each participant was randomly assigned to one of two conditions—the standard dictator game (outcome certainty) or hidden information game (outcome uncertainty) (see Fig. 2d,e). The standard dictator game (Fig. 2d) mirrored the baseline game in ref. ¹. Here, participants were presented with two options: A and B. Option A—the self-serving option—meant that the decider would receive 100 cents and the receiver would get 10 cents. Option B—the prosocial option—meant that the decider would get 85 cents and the receiver would get 85 cents. The hidden information game was adapted from the uncertainty treatment in ref. ¹. Participants first saw a game table that only specified the outcomes for the decider (that is, themselves), but not for the receiver. If choosing option A, participants would get 100 cents. If choosing option B, participants would get 85 cents. The receiver's outcome would depend on a virtual coin-flip, determining whether game 1 or game 2 would be played. In game 1, option A would result in 10 cents for the receiver, but in game 2 it would result in 85 cents. In game 1, option B would result in 85 cents for the receiver, but in game 2 it would result in 10 cents. Participants learned that they could reveal which game was played before making their decision at no cost. Hence, participants could choose option A (the self-serving option), option B or reveal which game was being played before making their decision. If participants chose to reveal the game, they saw which game was played and were then prompted to choose between option A and B.

To manipulate impact uncertainty, each participant except those in the control condition read the following description of the receivers: “MTurkers come from all walks of life, with different educational backgrounds and income levels. Some MTurkers, for instance, rank near the bottom of the income scale, and are very dependent on the money they earn on MTurk to help supplement their income to pay for food and shelter. Others rank in the middle-to-high end of the income scale, and use MTurk as a way to earn a bit of extra spending money; for example, to use for entertainment.” Participants were then randomly assigned to one of four

conditions: impact uncertainty, certain poor, certain rich or control. Participants in the impact uncertainty condition read that “The receiver might rank near the top of the income scale, or they might rank near the bottom, or somewhere in the middle.” Participants in the certain-poor condition read that: “The receiver ranks near the bottom of the income scale.” Participants in the certain-rich condition learned that: “The receiver ranks near the top of the income scale.” Participants in the control condition did not receive any information about the receiver.

Study 2. Participants. A total of 1,320 participants were recruited via Prolific—a crowdsourcing platform used to recruit participants online (prolific.ac) that is similar to MTurk. Sample size was determined using effect-size estimates from study 1 and aimed at replicating findings with a power of 0.80.

Procedure. The procedure was the same as in study 1, but we manipulated impact uncertainty differently to control for participants' beliefs about receivers' income. This time, participants in the uncertain condition read: “We pre-selected three receivers: a high-income receiver, a low-income receiver, and a middle-income receiver. You will be paired with one of these receivers at random. Thus, the receiver you are paired with might rank near the top of the income scale, or they might rank near the bottom, or somewhere in the middle.”

Study 3. Participants. A total of 742 participants were recruited via MTurk.

Procedure. To examine perceived social norms about the options presented in studies 1 and 2, we used a mixed design with impact uncertainty as the within-subjects factor (certain poor, certain rich or impact uncertainty) and outcome uncertainty (standard dictator game or hidden information game) as the between-subjects factor. All participants played an incentivized coordination game²¹. In this game, they were presented with the same instructions that participants saw in studies 1 and 2, but instead of deciding themselves between options A and B, they were asked to indicate how ‘socially appropriate’ or ‘socially inappropriate’ each of these options were (for details see Supplementary Methods).

Study 4. Participants. A total of 807 participants were recruited via Prolific.

Procedure. We used a between-subjects design with an independent variable of four levels (high-impact uncertainty, low-impact uncertainty, certain poor and control). The procedure was the same as in study 1 with the standard dictator game, but this time all participants—including those in the control condition—read the following general information about receivers: “Prolific workers come from all walks of life, with different educational backgrounds and income levels. Some Prolific workers, for instance, are poor (meaning their income is below the poverty threshold) and are very dependent on the money they earn on Prolific to help supplement their income to pay for food and shelter. Other Prolific workers have a high income, and use Prolific as a way to earn a bit of extra spending money; for example, to use for entertainment”. Deciders in the high-impact uncertainty condition were then told that there was a 50% chance that their receiver was poor, and a 50% chance that they were rich. In studies 1–3, we had told deciders that their receiver may be poor, rich or somewhere in the middle with a 33% probability each. We now used only the two extremes (that is, poor and rich) with a 50% split, because it is not intuitive what the norm for behaviour towards a middle-income receiver should be and, in fact, this aspect is not relevant to our research question. Participants in the low-impact uncertainty condition were told that there was a 10% chance their receiver was poor and a 90% chance their receiver was rich. The certain-poor condition was the same as in study 1. After participants made their decision in the dictator game, they answered questions about their demographics and completed individual differences measures. These included a measure of cognitive and affective empathy with well-established psychometric properties—the Questionnaire of Cognitive and Affective Empathy⁴²—and a measure of wise reasoning⁴³.

Study 5. Participants. A total of 862 participants were recruited via MTurk.

Procedure. In a between-subjects design, participants learned that we pre-selected three types of receivers—an extroverted receiver, an introverted receiver and a receiver who ranks in the middle—and that they would be randomly paired with one of them. Mirroring the impact uncertainty manipulation used in studies 1 and 2, participants in the certain-extrovert condition learned that the receiver was extroverted, participants in the certain-introvert condition learned that the receiver was introverted and participants in the control condition did not receive any information about the receiver. Thereafter, all participants played the standard dictator game (Fig. 1d).

Study 6. Participants. A total of 903 participants were recruited via MTurk.

Procedure. We used a three (uncertainty: impact uncertainty versus worst-case certainty versus control) by two (possibility: implicit versus explicit) between-subjects design to replicate our previous finding's robustness using a scenario-based paradigm and investigate whether impact uncertainty shifts people's representation of possible outcomes for others towards the worst-case possibility.

Manipulations for the uncertainty conditions were based on a fictive scenario set in the context of infectious disease. The implicit versus explicit possibility manipulation was based on a paradigm introduced recently³⁴ (see Supplementary Notes). The introductory text for the infectious disease scenario was the same across all conditions and read “Eight days after you arrived back from a lovely Safari trip to Tanzania, Africa, you feel unwell: feverish and dizzy. You go to the doctor and learn that you have the African flu. The doctor warns you that African flu is contagious: people you come into contact with may get infected (This middle part differed across conditions; see below). However, you still feel able to work and you really want to go to the office for finishing a project that is important for your career”. The scenario middle part differed across uncertainty conditions. In the impact uncertainty condition, participants learned there “is a chance that you would infect co-workers who are healthy people for whom the African flu is unproblematic (for example, a young person) so that they would only barely suffer. But there is also a chance that you would infect co-workers who are vulnerable people (for example, an old person) for whom the African flu is very dangerous so that they would suffer severely”. In the worst-case certainty condition, participants learned that most co-workers were vulnerable and hence that if they went to work, they were most likely to infect a vulnerable person. Participants in the control condition did not receive any information related to the vulnerability of their co-workers. Participants then made two possibility judgements (that is, the possibility that co-workers are vulnerable and the possibility of infecting co-workers) presented in random order either under time pressure (implicit condition) or without a time limit (explicit condition; see Supplementary Methods for instructions). Following their possibility judgement, participants proceeded to indicate whether or not they would go to work in the scenario they had read, on a seven-point Likert scale from ‘definitely not’ to ‘definitely’.

Study 7. Participants. A total of 468 participants were recruited via Prolific.

Procedure. We used a between-subjects design. In the low-salience condition, participants read that we had pre-selected one high-income, one middle-income and one low-income receiver, and that they would be randomly paired with one of these three receivers. In the high-salience condition, participants were additionally told: “The receiver might rank near the top of the income scale, or they might rank near the bottom, or somewhere in the middle.” In the control condition, participants did not receive any information about the receiver’s income before making their decisions. Thereafter, all participants played the standard dictator game (Fig. 1d).

Reporting Summary. Further information on experimental design is available in the Nature Research Reporting Summary linked to this article.

Code availability. The computer code that supports the findings of this study is available from the corresponding authors upon reasonable request.

Data availability. The data that support the findings of this study are available from the corresponding authors upon reasonable request.

Received: 26 April 2017; Accepted: 1 June 2018;

Published online: 09 July 2018

References

- Dana, J., Weber, R. A. & Kuang, J. X. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Econ. Theory* **33**, 67–80 (2007).
- Exley, C. L. Excusing selfishness in charitable giving: the role of risk. *Rev. Econ. Stud.* **83**, 587–628 (2015).
- Larson, T. & Capra, C. M. Exploiting moral wiggle room: illusory preference for fairness? A comment. *Judgm. Decis. Mak.* **4**, 467–474 (2009).
- Hamman, J. R., Loewenstein, G. & Weber, R. A. Self-interest through delegation: an additional rationale for the principal-agent relationship. *Am. Econ. Rev.* **100**, 1826–1846 (2010).
- Brock, J. M., Lange, A. & Ozbay, E. Y. Dictating the risk: experimental evidence on giving in risky environments. *Am. Econ. Rev.* **103**, 415–437 (2013).
- Feiler, L. Testing models of information avoidance with binary choice dictator games. *J. Econ. Psychol.* **45**, 253–267 (2014).
- Güth, W., Levati, M. V. & Ploner, M. On the social dimension of time and risk preferences: an experimental study. *Econ. Inq.* **46**, 261–272 (2008).
- Brennan, G., González, L. G., Güth, W. & Levati, M. V. Attitudes toward private and collective risk in individual and strategic choice situations. *J. Econ. Behav. Organ.* **67**, 253–262 (2008).
- Andreoni, J. & Bernheim, B. D. Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. *Econometrica* **77**, 1607–1636 (2009).
- Budescu, D. V., Rapoport, A. & Suleiman, R. Resource dilemmas with environmental uncertainty and asymmetric players. *Eur. J. Soc. Psychol.* **20**, 475–487 (1990).
- Gustafsson, M., Biel, A. & Gärling, T. Overharvesting of resources of unknown size. *Acta Psychol. (Amst.)* **103**, 47–64 (1999).
- De Kwaadsteniet, E. W., van Dijk, E., Wit, A., De Cremer, D. & de Rooij, M. Justifying decisions in social dilemmas: justification pressures and tacit coordination under environmental uncertainty. *Pers. Soc. Psychol. Bull.* **33**, 1648–1660 (2007).
- Shalvi, S., Gino, F., Barkan, R. & Ayal, S. Self-serving justifications: doing wrong and feeling moral. *Curr. Dir. Psychol. Sci.* **24**, 125–130 (2015).
- Nagel, T. What is it like to be a bat? *Phil. Rev.* **83**, 435–450 (1974).
- Harsanyi, J. C. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *J. Polit. Econ.* **63**, 309–321 (1955).
- Hammond, P. J. in *Interpersonal Comparisons of Well-Being* (eds Elster, J. & Roemer, J. E.) 200–254 (Cambridge Univ. Press, Cambridge, 1991).
- Falk, A. & Tirole, J. *Narratives, Imperatives and Moral Reasoning* Working Paper (mimeo Toulouse School of Economics, 2016); <https://pdfs.semanticscholar.org/8a0f/87ad25a8a51c7637b32b4e2fc5080f3963c4.pdf>
- Mazar, N., Amir, O. & Ariely, D. The dishonesty of honest people: a theory of self-concept maintenance. *J. Mark. Res.* **45**, 633–644 (2008).
- Haisley, E. C. & Weber, R. A. Self-serving interpretations of ambiguity in other-regarding behavior. *Games Econ. Behav.* **68**, 614–625 (2010).
- Van Dijk, E., Wit, A., Wilke, H. & Budescu, D. V. in *Contemporary Psychological Research on Social Dilemmas* (eds Suleiman, R., Budescu, D. V., Fischer, I. & Messick, D. M.) 315–331 (Cambridge Univ. Press, New York, NY, 2004).
- Krupka, E. L. & Weber, R. A. Identifying social norms using coordination games: why does dictator game sharing vary? *J. Eur. Econ. Assoc.* **11**, 495–524 (2013).
- Platt, M. L. & Huettel, S. A. Risky business: the neuroeconomics of decision making under uncertainty. *Nat. Neurosci.* **11**, 398–403 (2008).
- Buckholtz, J. W. et al. From blame to punishment: disrupting prefrontal cortex activity reveals norm enforcement mechanisms. *Neuron* **87**, 1369–1380 (2015).
- Savage, L. J. *The Foundations of Statistics* (Dover Publications, New York, NY, 1972).
- Davis, H. L., Hoch, S. J. & Ragsdale, E. K. E. An anchoring and adjustment model of spousal predictions. *J. Consum. Res.* **13**, 25–37 (1986).
- Greenstein, M. J. & Xu, X. My mug is valuable, but my partner’s is even more so: economic decisions for close others. *Basic Appl. Soc. Psychol.* **37**, 174–187 (2015).
- Hsee, C. K. & Weber, E. U. A fundamental prediction error: self–others discrepancies in risk preference. *J. Exp. Psychol. Gen.* **126**, 45–53 (1997).
- Ziegler, F. V. & Tunney, R. J. Who’s been framed? Framing effects are reduced in financial gambles made for others. *BMC Psychol.* **3**, 9 (2015).
- Andersson, O., Holm, H. J., Tyran, J.-R. & Wengström, E. *Risking Other People’s Money: Experimental Evidence on Bonus Schemes, Competition, and Altruism* IFN Working Paper 989 (IFN, 2013).
- Beisswanger, A. H., Stone, E. R., Hupp, J. M. & Allgaier, L. Risk taking in relationships: differences in deciding for oneself versus for a friend. *Basic Appl. Soc. Psychol.* **25**, 121–135 (2003).
- Zikmund-Fisher, B. J., Sarr, B., Fagerlin, A. & Ubel, P. A. A matter of perspective: choosing for others differs from choosing for yourself in making treatment decisions. *J. Gen. Intern. Med.* **21**, 618–622 (2006).
- Green, J. D., Burnette, J. L. & Davis, J. L. Third-party forgiveness: (not) forgiving your close other’s betrayer. *Pers. Soc. Psychol. Bull.* **34**, 407–418 (2008).
- Shalowitz, D. I., Garrett-Mayer, E. & Wendler, D. The accuracy of surrogate decision makers: a systematic review. *Arch. Intern. Med.* **166**, 493–497 (2006).
- Kelly, B., Rid, A. & Wendler, D. Systematic review: individuals’ goals for surrogate decision-making. *J. Am. Geriatr. Soc.* **60**, 884–895 (2012).
- Majesko, A., Hong, S. Y., Weissfeld, L. & White, D. B. Identifying family members who may struggle in the role of surrogate decision maker. *Crit. Care Med.* **40**, 2281–2286 (2012).
- Crockett, M. J., Kurth-Nelson, Z., Siegel, J. Z., Dayan, P. & Dolan, R. J. Harm to others outweighs harm to self in moral decision making. *Proc. Natl Acad. Sci. USA* **111**, 17320–17325 (2014).
- Engel, C. Dictator games: a meta study. *Exp. Econ.* **14**, 583–610 (2011).
- Bartling, B. & Fischbacher, U. Shifting the blame: on delegation and responsibility. *Rev. Econ. Stud.* **79**, 67–87 (2011).
- Aguilar, F., Brañas-Garza, P. & Miller, L. M. Moral distance in dictator games. *Judgm. Decis. Mak.* **3**, 344–354 (2008).
- Brañas-Garza, P. Poverty in dictator games: awakening solidarity. *J. Econ. Behav. Organ.* **60**, 306–320 (2006).
- Fong, C. & Oberholzer-Gee, F. Truth in giving: experimental evidence on the welfare effects of informed giving to the poor. *J. Public Econ.* **95**, 436–444 (2011).
- Reniers, R. L. E. P., Corcoran, R., Drake, R., Shryane, N. M. & Völlm, B. A. The QCAE: a Questionnaire of Cognitive and Affective Empathy. *J. Pers. Assess.* **93**, 84–95 (2011).
- Grossmann, I., Brienza, J. P. & Bobocel, D. R. Wise deliberation sustains cooperation. *Nat. Hum. Behav.* **1**, 0061 (2017).

44. Kogut, T. & Ritov, I. The “identified victim” effect: an identified group, or just a single individual? *J. Behav. Decis. Mak.* **18**, 157–167 (2005).
45. Small, D. A. & Loewenstein, G. Helping a victim or helping the victim: altruism and identifiability. *J. Risk Uncertain.* **26**, 5–16 (2003).
46. Small, D. A., Loewenstein, G. & Slovic, P. Sympathy and callousness: the impact of deliberative thought on donations to identifiable and statistical victims. *Organ. Behav. Hum. Decis. Process.* **102**, 143–153 (2007).
47. Ecker, U. K. H., Hogan, J. L. & Lewandowsky, S. Reminders and repetition of misinformation: helping or hindering its retraction? *J. Appl. Res. Mem. Cogn.* **6**, 185–192 (2017).
48. Putnam, A. L., Wahlheim, C. N. & Jacoby, L. L. Memory for flip-flopping: detection and recollection of political contradictions. *Mem. Cogn.* **42**, 1198–1210 (2014).
49. Stadler, M., Scharrer, L., Brummernhenrich, B. & Bromme, R. Dealing with uncertainty: readers’ memory for and use of conflicting information from science texts as function of presentation format and source expertise. *Cogn. Instr.* **31**, 130–150 (2013).
50. Carlyle, R. C. What are the chances of a devastating pandemic occurring in the next 50 years? *Huffington Post* https://www.huffingtonpost.com/quora/what-are-the-chances-of-a_b_3839785.html?guccounter=1 (2013).
51. Faul, F., Erdfelder, E., Buchner, A. & Lang, A.-G. Statistical power analyses using G*Power 3.1: tests for correlation and regression analyses. *Behav. Res. Methods* **41**, 1149–1160 (2009).
52. Buhrmester, M., Kwang, T. & Gosling, S. D. Amazon’s Mechanical Turk: a new source of inexpensive, yet high-quality, data? *Perspect. Psychol. Sci.* **6**, 3–5 (2011).
53. Paolacci, G., Chandler, J. & Ipeirotis, P. G. Running experiments on Amazon Mechanical Turk. *Judgm. Decis. Mak.* **5**, 411–419 (2010).
54. Phillips, J. & Cushman, F. Morality constrains the default representation of what is possible. *Proc. Natl Acad. Sci. USA* **114**, 4649–4654 (2017).

Acknowledgements

This work was supported by a Wellcome Trust Institutional Strategic Support Fund grant (204826/Z/16/Z) awarded to J.S., M.J.C., N.S.F. and G.K., a John Fell Fund award to M.J.C., and by the Oxford Martin Programme for Collective Responsibility for Infectious Disease. The work of J.S. for this paper was funded by a Wellcome Trust grant (104848/Z/14/Z). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The authors thank W. Sinnott-Armstrong, D. Batson, and members of the Crockett and MAD laboratories for helpful feedback.

Author contributions

A.K. and M.J.C. developed the research concept. A.K., M.J.C. and A.-M.N. designed the studies. Testing and data collection were performed by A.K. and A.-M.N. A.K. performed the data analysis and interpretation in collaboration with A.-M.N. and M.J.C. A.K., A.-M.N. and M.J.C. drafted the manuscript and all other authors provided critical revisions. All authors approved the final version of the manuscript for submission.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41562-018-0372-x>.

Reprints and permissions information is available at www.nature.com/reprints.

Correspondence and requests for materials should be addressed to A.K. or M.J.C.

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistical parameters

When statistical analyses are reported, confirm that the following items are present in the relevant location (e.g. figure legend, table legend, main text, or Methods section).

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- An indication of whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistics including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated
- Clearly defined error bars
State explicitly what error bars represent (e.g. SD, SE, CI)

Our web collection on [statistics for biologists](#) may be useful.

Software and code

Policy information about [availability of computer code](#)

Data collection

Data was collected via Amazon's Mechanical Turk and Academic Prolific. Surveys were implemented in Qualtrics.

Data analysis

Data was analyzed using SPSS, including the PROCESS macro that can be downloaded here: <http://www.processmacro.org/download.html>

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers upon request. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

All relevant data and SPSS code are available from the corresponding authors.

Field-specific reporting

Please select the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/authors/policies/ReportingSummary-flat.pdf](https://www.nature.com/authors/policies/ReportingSummary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	Our manuscript consists of 7 quantitative studies, of which Studies 1,2,4-7 are fully between-subject designs. Study 3 is a mixed within- and between-subjects design.
Research sample	All studies included male and female American adults (>18 years of age), residing within the United States. They were recruited via the online platforms AMT and Prolific, allowing for a more representative sample compared to university samples.
Sampling strategy	All studies aimed to achieve a minimum power of .80 and were based on effect size estimates drawn from previous studies. For details on power calculations see the methods section of the supplementary information.
Data collection	Data collection for all studies was conducted online, with participants recruited from Amazon's Mechanical Turk and Academic Prolific. Investigators were blinded to group allocation during data collection.
Timing	Studies 1-3 were collected in January and February 2015. Study 4 was collected in May 2017. Study 5 was collected in December 2015. Study 6 was collected in May 2017. Study 7 was collected in April 2016. The "Definitions Study" was collected in November 2017.
Data exclusions	For Study 4, participants who did not pass the attention check were excluded. For Study 6, participants who performed below a 60% accuracy during a training phase were excluded. These criteria were pre-established. See supplementary information methods part for details.
Non-participation	Some participants did not finish our online studies. However, we compared attrition rates across conditions and found these did not differ across conditions. See supplementary information methods part for details.
Randomization	Participants were randomized within our chosen research platform, Qualtrics. Thus, experimenters were unaware of which participants were allocated to which experimental condition. Randomization was not associated with any features of the participant, such as demographic or other individual differences variables.

Reporting for specific materials, systems and methods

Materials & experimental systems

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Unique biological materials
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input type="checkbox"/>	<input checked="" type="checkbox"/> Human research participants

Methods

n/a	Involvement in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics

Recruitment