Supplementary Methods

Power calculations

For Study 1, we used effect size estimates from related work by Dana and colleagues[1], and determined that a sample size of 800 participants would be required to achieve a power of .90. For Study 2, we used effect size estimates from our Study 1, and determined that we would need a sample of 1300 to replicate our findings with a power of .80. For Study 3, we used effect size estimates from Krupka and Weber[2] and determined that we would need a sample of 750 participants to achieve a power of .80. For Study 4, we sampled 800 participants based on the expectation that a sample of 200 participants per condition allow us to achieve the necessary power to detect small to medium (d = .3). Post-hoc power calculations confirmed that we achieved a power equal or greater than .84 in our logistic regression models. For Study 5, we used average effect size estimates from Study 1 and 2, and determined that we would need 850 participants to achieve a power of .90. We assumed a small effect size (d = .2) for the scenario-based Study 6 and determined we would need a sample of 900 to achieve a power of .80. Finally, for Study 7, we used we used average effect size estimates from Study 1 and 2, and determined that we would need a sample of 450 participants for a power of .80.

Attrition Rates

A recent paper[3] showed that online studies might violate the assumption of random assignment: sometimes participants drop out of one condition more often than from another. Since our conditions were brief, and did not differ in length, it appears unlikely that this would have affected our study. In Study 1, we had an attrition rate of 1.8%, in Study 2 of 1.4%, in Study 3 of 2.9%, in Study 4 of 1.7%, in Study 5 of 0.6%, in Study 6 of 3.2% and in Study 7 of 1.5%. Hence, attrition rates cannot explain the reported differences between conditions in our studies.

Definition Study

We wanted to test whether a naïve audience can easily distinguish between the two types of uncertainty in realistic examples. For this, we recruited a total of 70 participants were recruited via Amazon's Mechanical Turk (AMT). Participants first were presented with the definitions of both types of uncertainty used in the revised manuscript. The definitions were presented on separate pages. The first definition referring to outcome uncertainty read that this type is about "not knowing whether or not a negative outcome will occur", while the definition for impact uncertainty read that here uncertainty is about "not knowing how a negative outcome will impact another person". Both definitions were followed by two examples including the flu scenario (for outcome uncertainty: “a sick worker might wonder whether or not he will infect a co-worker if he goes to work”; impact uncertainty: “a worker might wonder how bad an infection would be for their co-worker”) and a graphic illustration of it.

On the subsequent page, participants had to answer a simple comprehension check about the respective types of uncertainty. Thereafter they were presented with an outcome and impact uncertainty case of three different real-world scenarios (yielding six versions) in a randomized order. Participants were asked to imagine themselves in these scenarios. One scenario asked them to imagine they are running a company and deciding whether to fire an employee (outcome uncertainty “You are uncertain about whether or not this employee will find a new job soon.”; impact uncertainty: “You are uncertain about how bad it will be for this employee if they do not find a new job soon.”). Another scenario asked them to imagine themselves as a tour guide deciding whether to take tourists to a volcanic island, for which scientists issued an eruption warning (outcome uncertainty: “You are uncertain about whether or not the volcano will erupt.”; impact uncertainty: “You are uncertain about how badly a volcano eruption would injure any tourists on your excursion.”). A third scenario asked them to imagine deciding whether to send a present for one’s best friend’s birthday at the last minute (outcome uncertainty: “You are uncertain about whether or not your present will arrive on time.”; impact uncertainty: “You are uncertain about how bad it would be for your friend if the present would not arrive on time.”). For each scenario, participants were then asked to indicate it was example of outcome or impact uncertainty, or whether they were not sure.

We found that the three scenarios exemplifying outcome uncertainty were correctly identified as such by 98 (firing an employee), 97 (leading tourists to volcano) and 94 (sending present to friend) percent of participants. The three scenarios exemplifying impact uncertainty were correctly identified as such by 97 (firing an employee), 94 (leading tourists to volcano), and 94 (sending present to friend) percent of participants. Only two participants indicated they were not sure for two of the scenarios. These results suggest that the difference between outcome and impact uncertainty is sufficiently intuitive for a naïve audience to reliably distinguish between these two types across various real-world examples.

Study 3

Procedure. Participants were introduced to the Dictator Game setup used in Studies 1-2 and instructed that they would be asked to predict how the other participants in their experiment evaluated the social appropriateness of the different possible choices available to the Decider. Socially appropriate behavior was defined as “consistent with moral or proper social behavior,” “behavior that most people agree is the ‘correct’ or ‘ethical’ thing to do”[2].

Participants then worked on an unrelated example situation. Thereafter, they learned that, at the end of the study, we would randomly select one of the Decider's choices. If the participant had given the response most frequently given by the other participants, then the participant would receive an additional 30 cents.

Study 4

Procedure. Different to the impact uncertainty manipulation used in Studies 1-3, instructions for this condition in Study 4 informed Deciders that their Receiver could be poor, or rich, with a 50% probability each. We introduced this modification to provide another test for the robustness of impact uncertainty's effect, and because the binary case is more intuitive. Note that for the binary case, uncertainty is maximal at the 50/50 chance distribution, whereas it is maximal at the 33/33/33 distribution for the previously used poor/middle-income/rich segmentation.

Study 6

Procedure (implicit/explicit manipulation). The implicit versus explicit possibility manipulation was based on a paradigm introduced recently by Phillips and Cushman[4], in which participants have to judge whether a given even event is possible or not either under time pressure (implicit possibility judgement), or without time pressure (explicit possibility judgement). To ensure participants have internalized the required responses for possible (pressing the [f] key) versus impossible (pressing the [j] key) judgments, they went through a training phase of 20 trials during which the words "possible" and "impossible" appeared each 10 times in a random order and participants had to press the adequate key. Participants who performed below 60 percent accuracy on these training trials were excluded from all analyses. Exemplary instructions for the implicit condition read "Please answer as quickly and accurately as you possibly can. Remember that you will only have about 1 second to respond to each event". Conversely, participants in the explicit condition read "Please take your time and carefully reflect on these questions. Make sure you do not answer too quickly or carelessly".

Supplementary Notes

Additional results

Study 1

We examined the effects of outcome and impact uncertainty on prosocial decision-making, using modified Dictator Games. Consistent with previous studies using this paradigm, when looking at those participants who revealed the game and faced the same pay-off as participants in the Standard Dictator Game (Figure 2d, Game 1), we found that more than 90% chose the prosocial option, suggesting that revealing the pay-off structure was prosocially motivated. When examining the participants who revealed the game and faced a pay-off in which the self-serving option coincides with the prosocial option (Figure 2b, Game 2), we found that all participants chose the self-serving / prosocial option.

Our results suggest that participants under impact uncertainty (i.e., in the uncertain condition) erred on the side of caution and assumed the worst possible impact outcome for the Receiver. One might suspect that the income of the participants could have driven the differences in self-interest between the different impact uncertainty conditions. We tested this possibility in a Generalized Linear Model with decision as a binary dependent variable and the information conditions, participants' income level, and the interaction between conditions and participants' income level as independent variables. We did not find a significant interaction between information condition and participants income level ($\chi^2(3, N = 800) = 3.23$, $p = .36$), indicating that the effect of information condition did not depend on participants' income level. We then examined the main effects for the impact uncertainty manipulation and income on prosocial decisions. This analysis revealed that participants with higher incomes were more prosocial ($\chi^2(1, N = 800) = 4.53$, $p = .03$), but the main effect of impact uncertainty remained significant ($\chi^2(3, N = 800) = 33.39$, $p < .001$), when controlling for income. Impact uncertainty affected participants' choices regardless of their own income level.

An alternative explanation is that participants in the uncertain condition assumed that most other participants on AMT – the platform we used to recruit participants – are in fact in the lowest income bracket, making it appropriate for them to choose relatively prosocial. To rule out this possibility we asked a subset of participants ($N = 414$), at the very end of the study, to estimate how many out of 100 workers on AMT would be in the highest income bracket, middle bracket, and lowest income bracket. One average, these participants believed that about 10 workers would be in the highest income bracket (MiN = 0, Max = 72; SD = 13.32), that about 40 workers would be in the medium income bracket (MiN = 0, Max = 100; SD = 17.87), and that about 50 workers would be in the low income bracket (MiN = 0, Max = 100; SD = 22.87). We then repeated the same control analysis as for income level to rule out the possibility that the impact uncertainty effects were driven by participants' estimated proportion of AMT workers in the lowest income bracket. We did not find a significant interaction between impact uncertainty condition and the estimated proportion of AMT workers in the lowest income bracket ($\chi^2(3, N = 414) = 1.08$, $p = .78$). When subsequently examining the main effects, we found no significant effect of the estimated proportion of AMT workers in the lowest income bracket on self-serving choices ($\chi^2(1, N = 414) = 0.08$, $p = .77$), but the main effect of impact uncertainty on prosociality remained significant ($\chi^2(3, N = 414) = 18.25$, $p < .001$).

Similarly, for the estimated proportion of the AMT workers in the highest income bracket, we found neither a significant interaction effect with impact uncertainty (χ^2((3, $N = 414) = 1.08$, $p = .77$), nor main effect for the estimated proportion ($\chi^2(1, N = 414) = 0.13$, $p = .72$), but the main effect of impact uncertainty on prosociality remained significant ($\chi^2(3, N = 414) = 18.16$, $p < .001$) when controlling for the estimated proportion of AMT workers in the highest income bracket. Taken together, these results indicate that the observed effect of impact uncertainty on prosocial behavior was not merely a reflection of an assumption that most AMT workers are in the lowest income bracket.

Study 2

We examined the independent and interactive effects of outcome and impact uncertainty on prosocial decision-making in a Generalized Linear Model predicting decision type (self-serving or prosocial) with separate regressors for outcome uncertainty (type of Dictator Game: Standard, Hidden Information), impact uncertainty (Receiver information: uncertain, certain-poor, certain-rich, control), and the interaction between outcome and impact uncertainty. We found a main effect of outcome uncertainty ($\chi^2(1, N = 1193) = 129.078$, $p < .001$), a main effect of impact uncertainty ($\chi^2(3, N = 1193) = 28.34$, $p < .001$), but no interaction effect ($\chi^2(3, N = 1193) = 2.45$, $p = .48$).

We next examined the effect of outcome uncertainty on prosociality more closely. In line with previous research , under outcome uncertainty (Hidden Information Game) participants behaved less prosocial than under outcome certainty (Standard Dictator Game) ($\chi^2(3, N = 1193) = 133.92$, $p < .001$, Cramer's $V = .33$). About a third of participants in the Hidden Information Game chose the prosocial option compared to more than two thirds of participants in the Standard Dictator Game. These percentages resemble the findings of previous laboratory studies[1].

Then, we examined the effect of impact uncertainty on prosociality. We first confirmed that participants in both studies 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). Participants in the certain-rich condition were less prosocial than those in the certain-poor condition ($\chi^2(1, N = 591) = 21.87$, $p < .001$, Cramer's $V = .19$). To investigate whether this difference was driven by increased generosity toward low-income Receivers or by decreased generosity toward 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. Participants in the certain-poor condition were significantly more prosocial than those in the control condition ($\chi^2(1, N = 595) = 14.54$, $p < .001$, Cramer's $V = .16$). Meanwhile, participants in the certain-rich condition were not significantly less prosocial than those in the control condition, ($\chi^2(1, N = 596) = .79$, $p = .37$, Cramer's $V = .037$). This suggests that the difference in prosociality between the certain-rich and certain-poor conditions was driven by increased generosity toward low-income Receivers.

To test our main prediction that impact uncertainty increases prosocial behavior, we compared the uncertain condition with each of the other three conditions. Participants were significantly more prosocial in the impact uncertainty conditions relative to the certain-rich conditions ($\chi^2(1, N = 598) = 8.63$, $p = .004$, Cramer's $V = .12$). These results speak against a self-serving exploitation of impact uncertainty. In contrast, the proportion of prosocial choices in the uncertain conditions was not significantly different from that in the certain-poor condition ($\chi^2(1, N = 597) = 3.19$, $p = .07$, Cramer's $V = .06$). And finally, participants in the impact uncertainty condition were significantly more prosocial than in the control condition ($\chi^2(1, N = 602) = 4.26$, $p = .04$, Cramer's $V = .08$). These results suggest that participants in the uncertain condition erred on the side of caution rather than exploiting uncertainty about the income level of the Receiver.

Again, we did not find a significant interaction between information condition and participants income level ($\chi^2(3, N = 1193) = 3.03$, $p = .35$), indicating that the effect of information condition did not depend on participants' income level. We then examined the main effects for the impact uncertainty manipulation and income on prosocial decisions. This analysis revealed that participants' incomes did not affect the results ($\chi^2(1, N = 1193) = 0.70$, $p = .43$), and the main effect of impact uncertainty remained significant ($\chi^2(3, N = 1193) = 26.10$, $p < .001$), when controlling for income. Impact uncertainty affected participants' choices regardless of their own income level.

Study 3

We replicated the opposing effects of outcome and impact uncertainty on prosociality by examining their distinct effects on social norms, using an incentivized coordination game[2] and the Dictator Games used in Study 1 and 2. The effects of both types of uncertainty on social norms mirrored those on behavior (Supplementary Figure 1). Under outcome uncertainty, the self-serving choice was perceived as less socially appropriate when the Receiver's outcome was certain, relative to when it was uncertain, ($p < .001$, $\eta^2 partial = .32$). Furthermore, the prosocial choice was perceived as less socially appropriate when the Receiver's outcome was uncertain, relative to when it was certain, ($p < .001$, $\eta^2 partial = .06$). Under impact uncertainty, participants perceived the self-serving choice to be less socially appropriate in the uncertainty condition than in the certain-rich condition, ($p < .001$, $\eta^2 partial = .09$), but more socially appropriate in the uncertain condition than in the certain-poor condition, ($p < .001$, $\eta^2 partial = .04$). When examining the appropriateness evaluations of the prosocial choice, we found that participants perceived the prosocial choice to be more socially appropriate in the uncertain condition than in the certain-rich condition, ($p < .001$, $\eta^2 partial = .04$), and as equally socially appropriate in the uncertain condition as in the certain-poor condition, ($p = .27$, $\eta^2 partial = .002$).

Supplementary Figure 1. Social appropriateness ratings of the prosocial (a) and self-serving decision (b) in the Standard Dictator Game and the Hidden Information Game. Social appropriateness ratings of the prosocial (c) and self-serving decision (d) for three Receiver information conditions: uncertain, certain-poor, and certain-rich. Error bars represent standard errors. *$p < .05$, n.s. = not significant.

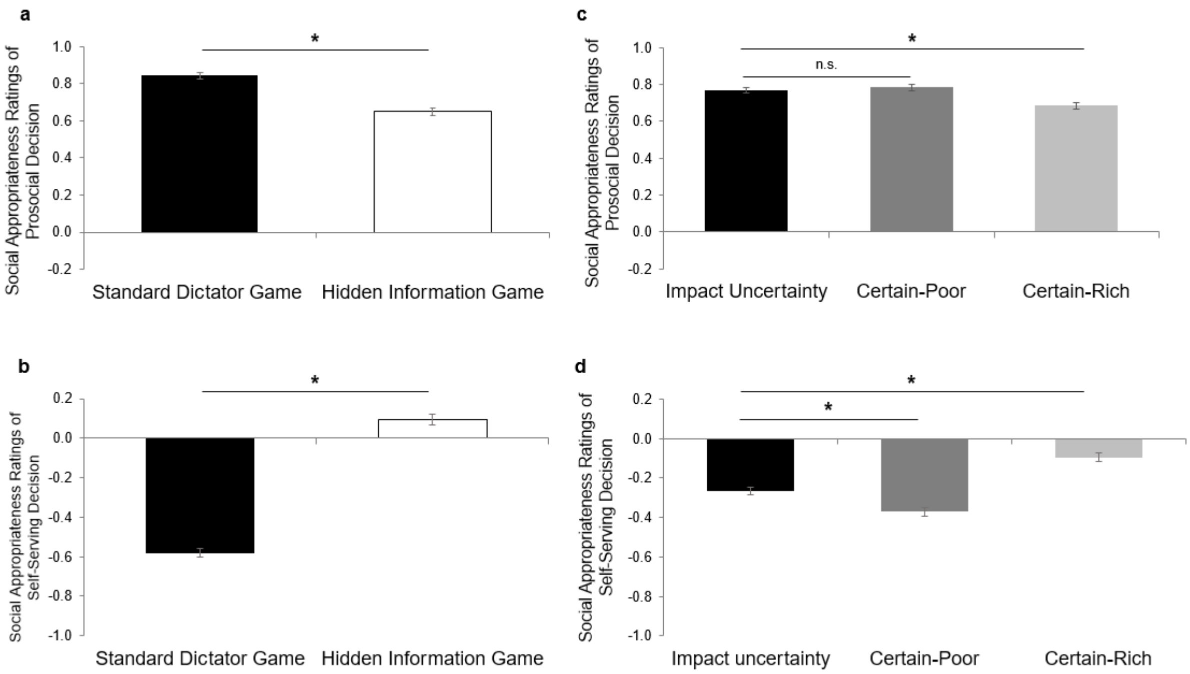

Study 4

Four participants failed the attention check and were excluded from all analyses, resulting in an overall sample of $N = 803$. Again, we did not observe a significant interaction between the conditional variable and participants' income on prosocial behavior ($p = .131$).

We next investigated the role of individual differences in empathy and wise reasoning for impact uncertainty's effect on prosociality. Mean scores on the cognitive and affective empathy subscales (Cronbach's alpha .915 and .833) differed significantly, $t(802) = 9.85$, $p < .001$, hence the subscales were not pooled. There were no interactions between either empathy subscale and the conditional variable (ps ≥ .481). Prosocial behavior was predicted by both cognitive ($\chi^2(1, N = 803) = 22.28$, $p < .001$) and affective empathy ($\chi^2(1, N = 803) = 14.45$, $p < .001$). Even though empathy is often conceptualized as trait[5], our manipulation might have affected participants' empathic mindset, which might then have mediated impact uncertainty's effect on prosociality. For instance, introducing impact uncertainty might give people a push towards considering others' perspectives, hence facilitating their ability to empathize with others. While there was no conditional effect on affective empathy ($p = .274$), cognitive empathy scores were indeed affected by the conditional manipulation, $F(3, 799) = 3.44$, $p = .016$, $\eta^2 = 0.03$. Bonferroni-corrected pairwise comparisons specified that this effect was driven by cognitive empathy scores being significantly higher under high impact uncertainty compared to the control condition, $p = .017$. Breaking cognitive empathy further down into its components[6], the conditional effect was significant only for the "perspective taking" component ($F(3, 799) = 3.98$, $p = .008$, $\eta^2 = 0.02$) but not the "online simulation" component ($p = .178$). To test the mediational hypothesis, we used Hayes' PROCESS tool in SPSS[7]. A bootstrap estimation with 5,000 samples[8] indicated the indirect effect was significant, $b = .02$, $SE = .01$, 95% CI = .001, .042. As the conditional effect on prosociality remained significant ($b = .20$, $SE = .09$, $p = .019$) after controlling for perspective taking as mediator ($b = .44$, $SE = .15$, $p = .004$), the results support a partial mediation of impact uncertainty's effect on prosociality via perspective taking.

Next, we tested whether wise reasoning also mediated the conditional effect on prosociality. Wise reasoning subscales were substantially correlated ($r \geq .460$, all $p < .001$) and thus we used their mean score for subsequent analyses[9]. Logistic regression analysis showed a significant main effect of wise reasoning on prosocial decisions ($b = .35$, $SE = .112$, $p = .003$), controlling for which the conditional effects of high impact uncertainty and certain poor remained significant ($b = .618$, $SE = .272$, $p = .023$; $b = .622$, $SE = .270$, $p = .021$). There was no interaction between the conditional and wise reasoning variables, $p = .311$. Neither the mean score, nor any of the wise reasoning subscales (range Cronbach's alpha = .756 to .863) were influenced by the conditional variable (mean score: $p = .227$; subscales: *p*s = .057 – $p = .887$; sum score: $p = .227$). When we tested the mediational pathway including the wise reasoning subscale that was the only one close to statistical significance for the conditional effect ("application of an outsider's vantage point", $p = .057$), the indirect effect was not significant ($b = .001$, $SE = .007$, 95% CI = -.011, .019).

Study 6

14 participants performed below the 60 percent accuracy cutoff and were excluded from all analyses, leaving a total sample of $N = 889$ ($n_{\text{impact uncertainty}} = 295$, $n_{\text{worst-case}} = 297$, $n_{\text{control}} = 297$). Since there was no interaction between the uncertainty and the explicit/implicit manipulations on prosocial intention ($p = .801$), we examined the two conditional effects separately. In addition to the conditional effect of uncertainty reported in our main manuscript, we found that participants who made implicit possibility judgments were less inclined to indicate they would go to work than those who made explicit possibility judgments, $U = 5.369$, $p = .021$, $\eta^2 = .749$. Logistic regression models showed there was also no interaction between the uncertainty and the explicit/implicit manipulations on possibility judgments (possibility to infect: $p = .226$; possibility of vulnerable co-worker: $p = .115$) so that we again separately examined the two conditional effects on the two possibility judgments. Note that 72 participants in the *implicit* conditions missed to make their possibility judgments on time, thus the following analyses are based on a total sample of $N = 817$. While the uncertainty manipulation did not affect possibility ratings about infecting a co-worker at all (*p*s ≥ .357), participants under *impact uncertainty* were significantly more inclined to judge it impossible to infect a vulnerable co-worker ($b = -1.010$, $SE = .402$, $p = .012$) and we observed the same trend for participants in the *worst-case certainty condition* ($b = -.706$, $SE = .419$, $p = .092$). For both possibility decisions (i.e., infection and vulnerability), participants

in the *implicit* conditions significantly more often judged the given outcome "impossible" compared to the *explicit* conditions (infect co-worker: $b = -.710$, $SE = .312$, $p = .023$; infect vulnerable co-worker: $b = -.880$, $SE = .307$, $p = .004$).

Study 7

We did not find a significant interaction between information condition and participants income level ($\chi^2(3, N = 466) = 1.19$, $p = .55$), indicating that the effect of information condition did not depend on participants' income level. We then examined the main effects for the impact uncertainty manipulation and income on prosocial decisions. This analysis revealed that participants' incomes did not affect the results ($\chi^2(1, N = 466) = 1.86$, $p = .17$), and the main effect of impact uncertainty remained significant ($\chi^2(3, N = 466) = 5.94$, $p = .05$), when controlling for income. Impact uncertainty affected participants' choices regardless of their own income level.

Effect Size Comparisons

We compared effect sizes across participants presented with low (10% chance of negative impact, Study 4), moderate (33% chance of negative impact, Studies 1 and 2), and high (50% chance of negative impact, Study 4) levels of impact uncertainty. The lowest effect size was observed for low uncertainty at Fisher's $Z_r = 0.05$, 95% CI [-0.04, 0.15]. Effect sizes for moderate impact uncertainty were larger, Fisher's $Z_r = 0.12$ and 0.08, 95% CIs [0.03, 0.22 and 0.01, 0.16], and similar to effect size observed for high impact uncertainty, Fisher's $Z_r = 0.12$, 95% CI [0.02, 0.22].

Supplementary References

1. 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).

2. 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).

3. Zhou, H. & Fishbach, A. The pitfall of experimenting on the web: How unattended selective attrition leads to surprising (yet false) research conclusions. *J. Pers. Soc. Psychol.* **111,** 493–504 (2016).

4. Phillips, J. & Cushman, F. Morality constrains the default representation of what is possible. *Proc. Natl. Acad. Sci.* **114,** 4649–4654 (2017).

5. Lamm, C., Batson, C. D. & Decety, J. The neural substrate of human empathy: Effects of perspective-taking and cognitive appraisal. *J. Cogn. Neurosci.* **19,** 42–58 (2007).

6. 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).

7. Hayes, A. F. *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. (Guilford Publications, 2017).

8. Shrout, P. E. & Bolger, N. Mediation in experimental and nonexperimental studies: new procedures and recommendations. *Psychol. Methods* **7,** 422 (2002).

9. Grossmann, I., Brienza, J. P. & Bobocel, D. R. Wise deliberation sustains cooperation. *Nat. Hum. Behav.* **1,** 0061 (2017).