



Human Cooperation When Acting Through Autonomous Machines

Celso M. de Melo^{a,1}, Stacy Marsella^b, and Jonathan Gratch^c

^aSensors and Electron Devices Directorate, US Army Research Laboratory, Playa Vista, CA 90094-2536; ^bCollege of Computer and Information Science, Northeastern University, Boston, MA 02115; and ^cInstitute for Creative Technologies, University of Southern California, Playa Vista, CA 90094-2536

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved January 7, 2019 (received for review October 15, 2018)

Recent times have seen an emergence of intelligent machines that act autonomously on our behalf, such as autonomous vehicles. Despite promises of increased efficiency, it is not clear whether this paradigm shift will change how we decide when our self-interest (e.g., comfort) is pitted against the collective interest (e.g., environment). Here we show that acting through machines changes the way people solve these social dilemmas and we present experimental evidence showing that participants program their autonomous vehicles to act more cooperatively than if they were driving themselves. We show that this happens because programming causes selfish short-term rewards to become less salient, leading to considerations of broader societal goals. We also show that the programmed behavior is influenced by past experience. Finally, we report evidence that the effect generalizes beyond the domain of autonomous vehicles. We discuss implications for designing autonomous machines that contribute to a more cooperative society.

autonomous vehicles | cooperation | social dilemmas

Recently, Google launched a service that allows a bot that sounds like a human to make calls and conduct business on people's behalf. This is just the latest example of autonomous intelligent technology that can act on our behalf. Recent times have seen increasing interest in such autonomous machines, including robots, personal home assistants, drones, and self-driving cars (1–4). These machines are changing the traditional ways we engage with others, and it is important to understand those changes. Recent research confirms that interacting through machines can affect the decisions we make with others (5). These findings showed that in ultimatum, impunity, and negotiation games, people were less likely to accept unfair offers if asked to program a machine to act on their behalf compared with direct interaction with proposers. This kind of experimental work has theoretical implications for our understanding of human decision making and practical implications for the design of autonomous machines. However, these effects are still not well understood and, in particular, it is unclear if cooperation is affected when people engage with others through an autonomous machine. Here we present experimental evidence that sheds light on this question and, at the same time, demonstrate the practical implications in the increasingly popular domain of autonomous vehicles.

Autonomous vehicles (AVs) promise to change the way we travel, move goods, and think of transportation in general (1). Not only are AVs supposed to be safer, many argue that they are an important step toward a carbon-free economy (6). Three reasons typically justify this claim: reduction in total number of traveled miles, reduction in traffic congestion, and the trend toward electric-powered AVs. However, this analysis mostly neglects cognitive factors that influence the way AVs will be programmed to drive. There has been some interest in moral dilemmas that AVs will face on the road, where decisions involving a cost in human life are made autonomously (2). However, less attention has been spent on the possibly more common case of social dilemmas that pit individual interests (e.g., comfort) against collective interests (e.g., the environment) that do not involve human life. Research shows that people can make distinct decisions in these dilemmas compared with situations involving sacred and moral values, such as when

human life is involved (7, 8). In social dilemmas, the highest reward occurs when the individual defects and everyone else cooperates; however, if everyone defects, then everyone is worse off than if they had all cooperated (9). Autonomous machines, including self-driving cars, are bound to face these kinds of dilemmas often, and thus we ask: Will people cooperate when engaging through these machines with others?

Theoretical Framework

When programming an autonomous machine, people are asked to think about the situation and decide ahead of time. This draws a parallel with the strategy method, often employed in experimental economics, whereby people are asked to complete decision profiles by reporting how they would decide for each possible eventuality in a task. A metareview of behavior in the ultimatum game revealed that, when asked to decide under the strategy method, participants were more likely to behave fairly compared with real-time interaction with their counterparts (10). This may occur because, on the one hand, people consider the counterparts' perspective (11) and, on the other hand, they rely on social norms, such as fairness, to provide a measure of consistency when considering the decisions for all possible outcomes (12). This research thus suggests that people should be more inclined to cooperate when programming an agent than when engaging in real-time direct interaction.

A compatible point of view is that programming a machine leads people to focus less on selfish short-term rewards, in contrast to what happens if the decisions are made on a moment-to-moment basis. In this sense, research has shown that the way a social dilemma is presented to the individual, or framed, can

Significance

Autonomous machines that act on our behalf—such as robots, drones, and autonomous vehicles—are quickly becoming a reality. These machines will face situations where individual interest conflicts with collective interest, and it is critical we understand if people will cooperate when acting through them. Here we show, in the increasingly popular domain of autonomous vehicles, that people program their vehicles to be more cooperative than they would if driving themselves. This happens because programming machines causes selfish short-term rewards to become less salient, and that encourages cooperation. Our results further indicate that personal experience influences how machines are programmed. Finally, we show that this effect generalizes beyond the domain of autonomous vehicles and we discuss theoretical and practical implications.

Author contributions: C.M.d.M., S.M., and J.G. designed research; C.M.d.M. performed research; C.M.d.M., S.M., and J.G. analyzed data; and C.M.d.M., S.M., and J.G. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence should be addressed. Email: celso.miguel.de.melo@gmail.com.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1817656116/-DCSupplemental.

Published online February 11, 2019.

impact the saliency of the selfish reward (13, 14). Complementing the view that social dilemmas involve a conflict between individual and collective interests, researchers have also proposed that social dilemmas present an inherent temporal conflict between a decision that leads to an immediate short-term reward (i.e., defection) and a decision that leads to a larger long-term reward (i.e., cooperation) (15, 16). When someone defects it is because people are temporally discounting distant rewards (17–19) and the short-term reward is perceived to be more salient. To avoid the temptation of the short-term reward, people can adopt a course of action that precludes the short-term reward (i.e., a precommitment) in favor of the higher long-term reward (20–22)—in our case, programming the machine accordingly. Moreover, once people make a commitment, they tend to behave in a manner that is consistent with that commitment (23, 24). Thus our hypothesis is that programming machines causes short-term rewards to become less salient, leading to increased cooperation.

This hypothesis is also in line with predictions coming from work on construal-level theory (25). According to this theory, when people mentally represent or construe a situation at a higher level, they focus on more abstract and global aspects of the situation; in contrast, when people construe a situation at a lower level, they focus on more specific context-dependent aspects of the situation. Building on this theory, Agerström and Björklund (26, 27) argued that, since moral principles are generally represented at a more abstract level than selfish motives, moral behaviors should be perceived as more important, with greater temporal distance from the moral dilemma [see Gong and Medin (28) for a dissenting view]. In separate studies, Agerström and Björklund showed that people made harsher moral judgments—in dilemmas such as the “recycling dilemma,” which never involved costs in human life—of others’ distant-future morally questionable behavior (26), and that people were more likely to commit to moral behavior when thinking about distant versus near-future events (27). In the fourth experiment in one of their studies (27), they further showed that the saliency of moral values mediated the effect of temporal distance on moral behavior. Kortenkamp and Moore (19) further showed that individuals with a chronic concern for an abstract level of construal—that is, those who were high in consideration for the future consequences of their behavior—exhibited higher levels of cooperation. In a negotiation setting, Henderson et al. (29) and De Dreu et al. (30) showed that individuals at a high construal level negotiated more mutually beneficial and integrative agreements. Finally, Fujita et al. (31) noted that high-level thinking can enhance people’s sense of control; increased self-control, in turn, can lead individuals to precommit to the more profitable distant reward (20–22). Therefore, the hypothesis is that programming machines leads to increased cooperation because it motivates more abstract thinking, which is likely to decrease the saliency of selfish short-term rewards.

Finally, building on this work on construal-level theory, Giacomantonio et al. (32) proposed that, rather than simply promoting cooperation, abstract thinking reinforces one’s values. Thus, under high-construal thinking, prosocials—as measured by a social orientation scale (33)—would be more likely to cooperate, whereas individuals with a selfish orientation, or proselves, would be more likely to defect. In support of this thesis, Giacomantonio et al. showed that, in the ultimatum game and negotiation, prosocials were more likely to cooperate under high construal but proselves were more likely to compete. This work suggests an alternative hypothesis, whereby prosocials would be more likely to program machines to cooperate compared with direct interaction; complementarily, proselves would be more likely to program machines to defect compared with direct interaction.

Overview of Experiments

To study people’s behavior when programming machines and the underlying mechanism, we ran five experiments where participants engaged in a social dilemma—the *n*-person prisoner’s dilemma—either as the owners of an autonomous vehicle or as the drivers themselves. Economic games such as the prisoner’s

dilemma are abstractions of prototypical real-life situations and are ideal for studying the underlying psychological mechanisms in controlled experiments (34). The decisions that participants made, moreover, always had real financial consequences, as the points they earned in the games were converted to tickets for lotteries—one per experiment—for a \$30 prize. Finally, to prevent any reputation concerns (35), the experiments we conducted were always fully anonymous—that is, the participants were anonymous to each other and to the experimenters (see *Methods* for details on how this was accomplished).

The prisoner’s dilemma was presented to the participants as an environmental social dilemma (16), where they would choose between two options: turning on the air conditioner (AC)—defection—which increased comfort at the expense of the environment because more gas was consumed, and turning off the AC—cooperation—which reduced comfort in favor of the environment because less gas was consumed. The critical comparison was between the cooperation rate—how often the AC was turned off—when acting through an AV vs. driving themselves.

When studying machines that act on behalf of humans, it is also important to clarify how much autonomy is given to these machines. At one extreme, the decisions made by the machine can be fully specified by humans; at the other extreme, the machine can make the decision by itself with minimal input from humans. The degree of autonomy is an important factor that influences the way people behave with machines (36). Effectively, research suggests that the degree of thought and intentionality behind a decision can impact people’s reactions (37). In this paper, we focus on machines that have minimal autonomy. Thus, our machines make decisions that are fully specified by the humans they represent. The rationale is that, if engaging with others through machines that have minimal autonomy affects social behavior, the effect is likely to exist (and possibly be stronger) with increased autonomy.

In experiment 1, we tested and confirmed that people cooperated more when programming their AVs compared with direct interaction with others. This effect was not moderated by the participants’ social value orientations. In experiment 2, we showed that programming caused selfish short-term rewards to become less salient, and this led to increased cooperation. Experiment 2b revealed that this effect was robust even when participants could reprogram their vehicles during the interaction. Experiment 3 showed that participants adjusted their programs based on the (cooperative or competitive) behavior they experienced from their counterparts. Finally, experiment 4 showed that the effect also occurred in an abstract social dilemma, suggesting that it generalizes beyond the domain of autonomous vehicles.

Experiment 1: Do People Program AVs to Cooperate More?

This experiment followed a simple two-level between-participants design: *autonomy* (programming vs. direct interaction). Participants engaged with three counterparts in the four-person prisoner’s dilemma. The payoff table for this game is shown in Fig. 1A (*top center*). Participants made simultaneous decisions without communicating with each other. For instance, if the participant chose to turn off the AC and the remaining players chose to turn it on, then the participant got 4 points whereas the others got 12 points (see the second column in the payoff table). If everybody chose to turn off the AC, everybody got 16 points (see the fifth column in the payoff table). However, if everyone decided to turn on their AC, then everyone got only 8 points (see the first column in the payoff table). Participants engaged in 10 rounds of this dilemma with the same counterparts. For this first experiment, we wanted to exclude any strategic considerations that could occur due to repeated interaction. To accomplish this, participants were informed that they would not learn about the decisions others made until the end of the last round, when they would find out how much they had earned in total. Before engaging in the task, participants were quizzed to make sure that they understood the instructions and were not able to proceed until they answered all questions correctly.




Fig. 1. Experimental conditions in experiment 2. (A) Participants drive the vehicle directly and are instructed to focus on the outcome for the self; the monetary reward is made salient. (B) Participants program the vehicle to act on their behalf and are instructed to focus on the outcome for all; the environment is made salient.

In the case of direct interaction, participants saw their car driving down a road and come to a stop; at that point, they were given the opportunity to make the decision on whether to turn on the AC. After making their decision, the car continued driving and the procedure was repeated for every round. In the autonomous vehicle case, participants programmed their car, before the first round, to make whichever decision they wanted in each round. Then a simulation started showing the car driving down a road and coming to a stop. At that point, without intervention from the participant, the car played the decision it was programmed to make autonomously. This procedure was repeated for every round. (Movie S1 shows the software used in every experiment and demonstrates all experimental conditions.)

Before starting the task, participants had to wait approximately 30 seconds for other participants to join; additionally, after submitting their decisions in each round, participants also had to wait for other participants to finish making their decisions. However, to increase experimental control, participants always engaged with a computer script that simulated the other participants. Similar experimental manipulations have been used in other experiments studying behavior involving intelligent machines (5, 38, 39). Participants were fully debriefed about this experimental procedure at the end of the experiment.

The main measure was cooperation rate, averaged across the 10 rounds. Before starting the task, we also collected the participants' social value orientation (SVO) using the slider scale (40). We recruited 98 participants from an online pool, Amazon Mechanical Turk. Previous research shows that studies performed on Mechanical Turk can yield high-quality data and successfully replicate the results of behavioral studies performed on traditional pools (41). Since some gender effects on cooperation have been reported in the literature (42, 43), we began by running a gender \times autonomy between-participants ANOVA; however, because we did not find any statistically significant main effects of gender or gender \times autonomy interactions in any of the experiments, we did not address this issue further (but see *SI Appendix* for a full analysis of gender). To compare the cooperation rate for programming vs. direct interaction, we ran an independent-samples t test. The results confirmed that participants were statistically significantly more likely to turn off the AC—that is, cooperate—when programming their AV to act on their behalf ($M = 0.64$, $SE = 0.06$) than when driving themselves ($M = 0.47$, $SE = 0.05$), $t(96) = 2.33$, $P = 0.022$, $r = 0.23$. To test if social value orientation was moderating this effect (32), we ran an SVO \times autonomy between-participants ANOVA, which revealed no significant SVO \times autonomy interaction, $F(1, 94) = 1.636$, $P = 0.202$. Thus, we found no evidence supporting the contention that prosocials would program AVs to cooperate more whereas proselves would program AVs to cooperate less. Similar patterns were found in all experiments and so SVO was not addressed further (but see *SI Appendix* for a complete analysis of SVO in all experiments).

Experiment 2: What Is the Mechanism?

Experiment 1 showed that people programmed AVs to cooperate more than when driving themselves. In experiment 2, we wanted to get insight on the mechanism behind this effect. As described in *Theoretical Framework*, we compared complementary mechanisms based on focus on self, saliency of the short-term reward, fairness, and high-construal reasoning. To accomplish this, we introduced a second experimental factor, *saliency*, that instructed the participant that turning on the AC was more likely than the other option to lead to a higher payoff for themselves (Fig. 1A); in the other case, in contrast, participants were informed that turning on the AC was less likely than the other option to be environmentally-friendly (Fig. 1B). To compare saliency with more traditional manipulations of social motivation (44), we introduced a third factor, *focus*, that simply instructed participants to focus on how the decision would “affect yourself” (Fig. 1A) vs. “yourself and the others” (Fig. 1B). Finally, we administered four subjective scales, after completing the task, that represented four possible mediators: focus on self (e.g., “I was mostly worried about myself”), short-term reward saliency (e.g., “I was focused on the monetary payoff”), fairness (e.g., “I wanted to be fair”), and high-construal reasoning (e.g., “I was thinking in more general terms, at a higher level, about the situation”). See *SI Appendix* for full details on these scales.

The experiment thus followed a $2 \times 2 \times 2$ between-participants factorial design: *autonomy* (programming vs. direct interaction) \times *saliency* (monetary reward vs. environment) \times *focus* (self vs. collective). The autonomy factor and the experimental procedure were similar to experiment 1, and we used an identical financial incentive. We recruited 334 participants from Amazon Mechanical Turk. To analyze cooperation rate, we ran an autonomy \times saliency \times focus between-participants ANOVA. The average cooperation rates are shown in Fig. 24. Once again, participants

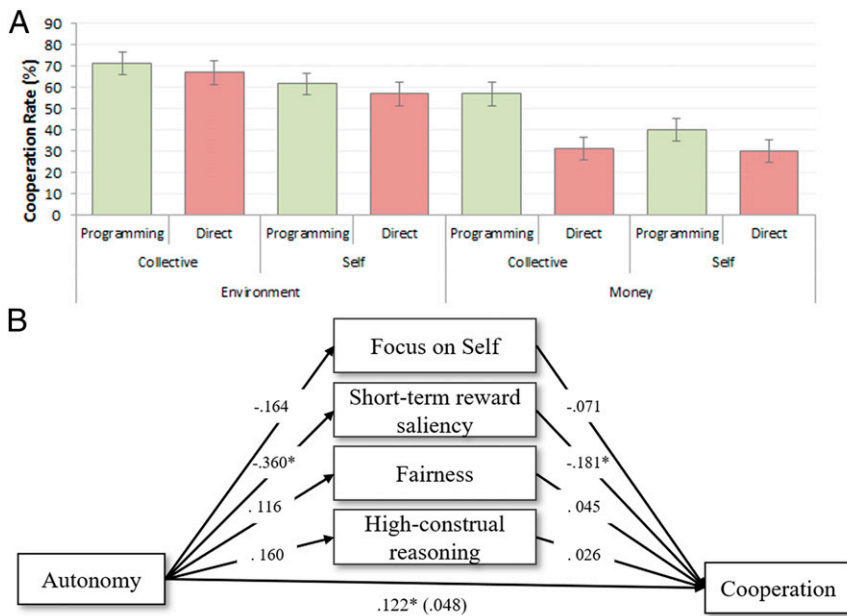


Fig. 2. Results in experiment 2. (A) Cooperation rates for all conditions. The error bars represent SEs. (B) Analysis of statistical mediation of focus on self, short-term reward saliency, fairness, and high-construal reasoning on the effect of autonomy on cooperation. * $P < 0.05$.

were more likely to turn off the AC—cooperate—when programming their AVs ($M = 0.58, SD = 0.03$) than when driving themselves ($M = 0.46, SE = 0.03$), $F(1, 326) = 8.56, P = 0.004$, partial $\eta^2 = 0.026$. In addition, participants cooperated more when instructed to focus on the collective ($M = 0.57, SE = 0.03$) than the self ($M = 0.47, SE = 0.03$), $F(1, 326) = 6.22, P = 0.013$, partial $\eta^2 = 0.019$. There was, however, no statistically significant autonomy \times focus interaction [$F(1, 326) = 1.02, ns$]. This suggests that, even though cooperation is impacted by experimentally manipulating social motivation, as has been shown in the past (44), this factor did not appear to explain the effect of autonomy. Finally, participants cooperated more when the environment was more salient ($M = 0.64, SE = 0.03$) than the monetary reward ($M = 0.40, SE = 0.03$), $F(1, 326) = 41.76, P < 0.01$, partial $\eta^2 = 0.114$. Moreover, in contrast to the previous case, this time there was a trend for an autonomy \times saliency interaction, $F(1, 326) = 2.95, P = 0.087$, partial $\eta^2 = 0.009$: people cooperated more when programming AVs than in direct interaction, but only when the monetary reward was salient.

To complement and reinforce this analysis, we ran a multiple-mediation analysis (45) where we tested if focus on self, short-term reward saliency, fairness, and high-construal reasoning mediated the effect of autonomy on cooperation. The independent variable, autonomy, was binary-coded: 0 for direct interaction; 1 for programming. The dependent variable was cooperation rate. Fig. 2B shows the mediation analysis results. (See *SI Appendix, Table S1*, for details on bootstrapping tests for statistical significance of the indirect effects.) The results indicated that the total indirect effect was statistically significant (0.073, $P = 0.04$) and, more important, the indirect effect for short-term reward saliency was the only one that was statistically significant (0.063, $P = 0.03$). That is, the corresponding 95% bootstrapping confidence interval did not include zero. The analysis also indicated that the total effect (0.122, $P = 0.003$) became statistically nonsignificant once we accounted for the effect of the mediators: The direct effect was statistically nonsignificant (0.048, ns). In sum, the mediation analysis suggested that the effect of autonomy on cooperation is proximally caused by reduced saliency of the short-term reward.

Experiment 2b: What if Reprogramming Were Permitted?

In the previous experiment, participants programmed their autonomous vehicles in the beginning and were never given the opportunity to reprogram during the task. In other words, participants were required to precommit to their choice at the start of the task. Research indicates that precommitment can be

effective in minimizing the temptation to choose the short-term reward (20–22). However, does the effect still occur if people are not required to precommit? The answer to this question matters because, in practice, manufacturers are unlikely to force customers to precommit their decisions and never give them the chance to change their decisions based on their driving experience. Experiment 2b thus replicated the previous experiment but with one important change: If they desired, participants could reprogram their vehicles at the start of each round.

To remove strategic considerations, the previous experiment had participants engage with the same counterparts throughout all rounds but at the same time prevented them from learning what the counterparts decided in each round. In this experiment, we wanted to use a more natural paradigm that still minimized strategic reasoning. In experiment 2b, first, participants engaged with different counterparts in each round. Second, participants learned, at the end of each round, about the counterparts' decisions. The decisions the counterparts made were neither too cooperative nor too competitive: in one half of the rounds, two counterparts defected (rounds 1, 3, 6, 8, and 10); in the other half, two counterparts cooperated (rounds 2, 4, 5, 7, and 9).

This experiment thus followed the same design and procedure as experiment 2, with three exceptions. First, participants in the programming condition were given the opportunity to reprogram the vehicle at the start of every round following the first round, where everyone was required to program the vehicle. When reprogramming, participants saw a screen with all of the current decisions chosen by default and were allowed to change the decisions for every round that had yet to occur. Second, participants were told that each round would be played with different counterparts and that they would never engage with the same counterpart for more than one round. Finally, participants were given information about the counterparts' decisions at the end of each round.

We recruited 351 participants from Amazon Mechanical Turk. Similarly to experiment 2, participants cooperated more when programming their AVs ($M = 0.54, SD = 0.02$) than when driving themselves ($M = 0.47, SE = 0.02$), $F(1, 343) = 4.59, P = 0.033$, partial $\eta^2 = 0.013$. Participants also cooperated more when the environment was more salient ($M = 0.63, SE = 0.02$) than the monetary reward ($M = 0.37, SE = 0.03$), $F(1, 343) = 56.88, P < 0.01$, partial $\eta^2 = 0.142$. There was also a replication of the trend for autonomy \times saliency interaction, $F(1, 343) = 2.96, P = 0.086$, partial $\eta^2 = 0.009$: people cooperated more when programming AVs than in direct interaction, but only when the monetary

reward was salient. In this experiment, however, there was no statistically significant effect of focus on the collective ($M = 0.51$, $SE = 0.03$) vs. the self ($M = 0.50$, $SE = 0.02$), $F(1, 343) = 0.720$, ns. The failure to replicate this effect suggests that the effect of focus was small in this setting compared with the effect of saliency. Overall, the experiment reinforced that people are likely to focus less on the short-term reward and cooperate more when programming their vehicles, even when given the opportunity to change their minds throughout the interaction (i.e., in the absence of a precommitment).

Experiment 3: What if Others Behaved Cooperatively or Competitively?

In the previous experiments, we created conditions to minimize strategic interactions that may occur due to repeated interaction; in this experiment, in contrast, we explicitly studied participants' behavior when they engaged with counterparts who behaved either cooperatively or competitively in repeated interaction. Based on recent work showing that people adjust their negotiation programs based on personal experience (46), we expected participants to adapt their programming behavior according to their experience in the repeated dilemma. Furthermore, we wanted to study if participants would differentiate their behavior when the interaction was with different counterparts in every round—repeated one-shot games—or always with the same counterparts—a single repeated game. This is relevant since, in practice, people are more likely to engage in their driving with different drivers and thus it is important to understand if their experience spills over to new drivers with whom there is no interaction history.

The experiment followed a $2 \times 2 \times 2$ between-participants factorial design: *autonomy* (programming vs. direct interaction) \times *counterpart behavior* (cooperative vs. competitive) \times *persistence* (always the same counterparts vs. different counterparts in every round). In the cooperative condition, in one half of the rounds, two counterparts cooperated (rounds 1, 3, 4, 7, and 8); in the other half, every counterpart cooperated (Rounds 2, 5, 6, 9, and 10). In the competitive condition, in one half of the rounds, two counterparts defected (rounds 1, 3, 4, 7, and 8); in the other half, every counterpart defected (rounds 2, 5, 6, 9, and 10). We recruited 339 participants from Amazon Mechanical Turk. The cooperation rate for each condition, across rounds, is shown in Fig. 3. To analyze the data we ran a between-participants ANOVA that confirmed a trend for a main effect of autonomy, $F(1, 331) = 2.64$, $P = 0.105$, partial $\eta^2 = 0.008$, with participants cooperating more when programming their AVs ($M = 0.54$, $SD = 0.02$) than when driving themselves ($M = 0.48$, $SE = 0.02$); this effect becomes statistically significant if we look only at the last round, $F(1, 331) = 5.17$, $P = 0.024$, partial $\eta^2 = 0.015$. The results confirmed an effect of counterpart strategy, with participants cooperating more with cooperative ($M = 0.57$, $SE = 0.02$) than with competitive counterparts ($M = 0.45$, $SE = 0.02$), $F(1, 331) = 12.96$, $P < 0.001$, partial $\eta^2 = 0.038$. However, there was no significant autonomy \times counterpart strategy interaction, $F(1, 331) = 1.30$, $P = 0.255$, which suggests that counterparts' behavior was just as likely to influence participants' programming as driving behavior. Somewhat

surprisingly, we found no effect of persistency, $F(1, 331) = 1.10$, $P = 0.295$, which suggests that people were still adjusting their behavior with new counterparts, with whom they had no history, based on their experience from previous rounds with different counterparts. A more detailed analysis of how cooperation rate unfolded across rounds is presented in the *SI Appendix*.

Experiment 4: Does the Effect Generalize?

This experiment sought to study whether the effect identified in the previous experiments generalizes beyond this environment preservation task involving autonomous vehicles. To test this we had participants engage in an abstract social dilemma, without any mention of AVs. Like experiment 1, this experiment followed a two-level between-participants design: *autonomy* (programming vs. direct interaction). The experimental procedure was similar to the one followed in experiment 2b, except that we did not consider the saliency and focus factors. The important difference was that the n -person social dilemma was presented abstractly and there was no mention of self-driving cars or environment. In this case, the participant made a choice between option A, corresponding to cooperation, and option B, corresponding to defection (*SI Appendix*, Fig. S1, shows a screenshot). The payoff matrix was exactly the same as in the previous experiments. To justify the programming condition, we mentioned that the purpose of the experiment was to study how people make decisions through "computer agents" that act on their behalf. Thus, in this experiment, instead of programming self-driving cars, participants programmed a computer agent. As in experiment 2b, participants were allowed to reprogram the agents at the start of each round. We recruited 103 participants from Amazon Mechanical Turk. To analyze cooperation rates, we ran an independent-samples t test for programming vs. direct interaction. In line with the previous experiments, the results showed that participants cooperated more when programming their agents ($M = 0.45$, $SE = 0.05$) than when engaging directly with others ($M = 0.31$, $SE = 0.04$), $t(101) = 2.21$, $P = 0.029$, $r = 0.22$.

Discussion

As autonomous machines that act on our behalf become immersed in society, it is important that we understand whether and, if so, how people change their behavior and what are the corresponding social implications. Here we show that this analysis is incomplete if we ignore the cognitive factors shaping the way people decide through these machines. In the increasingly relevant domain of autonomous vehicles, we showed that programming vehicles ahead of time led to an important change in the way social dilemmas are solved, with people being more cooperative—that is, more willing to sacrifice individual for collective interest—than if driving the vehicle directly. We also showed that this effect prevailed even if participants were allowed to reprogram their machines (experiment 2b) and independently of whether others behaved cooperatively or competitively (experiment 3). Finally, we showed that this effect was not limited to the domain of autonomous vehicles. In line with earlier work showing that people program machines to act more fairly in abstract games (5), participants still showed increased cooperation when programming a machine in an abstract version of the dilemma that had no mention of AVs (experiment 4).

We have provided insight into the mechanism driving the effect, emphasizing the role programming machines play in lowering the saliency of the short-term reward when deciding whether to cooperate (experiment 2). This is in line with research showing that introducing an opportunity to make a precommitment—in this case, by programming the machine ahead of time—can reduce the temptation to choose the short-term reward (20–22). However, the results also clarify that the effect still occurs in the absence of precommitment (experiment 2b), in that participants were motivated to maintain the commitment they made in the initial round (23, 24). The findings also align with research indicating that, when people are asked to report their decisions ahead of time, they tend to behave more fairly




Fig. 3. Cooperation rates across the 10 rounds in experiment 3.

than when making the decision in real time (10–12), and with prior work showing that cooperation can be encouraged by the way a social dilemma is framed (13, 14). Finally, these findings are compatible with earlier work by construal-level theorists suggesting that the adoption of high-construal abstract thinking—as is motivated by programing an autonomous machine—increases the saliency of moral values (19, 26, 27, 30) and self-control (31), which can encourage cooperation.

We focused on the case where the decision is made by the owner of the autonomous machine. This allowed us to identify an important effect and mechanism for cooperation that is made salient by virtue of having to program the machine. However, in practice this decision may be distributed across multiple stakeholders with competing interests, including government, manufacturers, and owners. In the case of a moral dilemma, Bonnefon et al. (2) showed that people prefer other people's AVs to maximize preservation of life (even if that means sacrificing the driver) but their own vehicle to maximize preservation of the driver's life. As these issues are debated, it is important that we understand that, in the possibly more prevalent case of social dilemmas—where individual interest is pitted against collective interest in situations that do not involve a cost in human life—autonomous machines have the potential to shape how the dilemmas are solved, and thus these stakeholders have an opportunity to promote a more cooperative society.

- Waldrop MM (2015) Autonomous vehicles: No drivers required. *Nature* 518:20–23.
- Bonnefon J-F, Shariff A, Rahwan I (2016) The social dilemma of autonomous vehicles. *Science* 352:1573–1576.
- Floreano D, Wood RJ (2015) Science, technology and the future of small autonomous drones. *Nature* 521:460–466.
- Stone R, Lavine M (2014) Robots. The social life of robots. Introduction. *Science* 346: 178–179.
- de Melo C, Marsella S, Gratch J (2017) Social decisions and fairness change when people's interests are represented by autonomous agents. *Auton Agents Multi Agent Syst* 32:163–187.
- Alexander-Kearns M, Peterson M, Cassidy A (2016) The impact of vehicle automation on carbon emissions. Retrieved from Center for American Progress Website. <https://www.americanprogress.org/issues/green/reports/2016/11/18/292588/the-impact-of-vehicle-automation-on-carbon-emissions-where-uncertainty-lies/>.
- Tetlock PE, Kristel OV, Elson SB, Green MC, Lerner JS (2000) The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *J Pers Soc Psychol* 78:853–870.
- Dehghani M, Carnevale P, Gratch J (2014) Interpersonal effects of expressed anger and sorrow in morally charged negotiation. *Judgm Decis Mak* 9:104–113.
- Dawes R (1980) Social dilemmas. *Annu Rev Psychol* 31:169–193.
- Oosterbeek H, Sloof R, Van de Kuilen G (2004) Cultural differences in ultimatum game experiments: Evidence from a meta-analysis. *Exp Econ* 7:171–188.
- Güth W, Tietz R (1990) Ultimatum bargaining behavior: A survey and comparison of experimental results. *J Econ Psychol* 11:417–449.
- Rauhut H, Winter F (2010) A sociological perspective on measuring social norms by means of strategy method experiments. *Soc Sci Res* 39:1181–1194.
- Pruitt D (1970) Motivational processes in the decomposed prisoner's dilemma game. *J Pers Soc Psychol* 14:227–238.
- Lieberman V, Samuels SM, Ross L (2004) The name of the game: Predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. *Pers Soc Psychol Bull* 30:1175–1185.
- Dewitte S, De Cremer D (2001) Self-control and cooperation: Different concepts, similar decisions? A question of the right perspective. *J Psychol* 135:133–153.
- Joireman J (2005) Environmental problems as social dilemmas: The temporal dimension. *Understanding Behavior in the Context of Time*, eds Strathman A, Joireman J (Lawrence Erlbaum Associates, London), pp 289–304.
- Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: A critical review. *J Econ Lit* 15:351–401.
- Mannix E (1990) Resource dilemmas and discount rates in decision making groups. *J Exp Soc Psychol* 27:379–391.
- Kortenkamp KV, Moore CF (2006) Time, uncertainty, and individual differences in decisions to cooperate in resource dilemmas. *Pers Soc Psychol Bull* 32:603–615.
- Wertenbroch K (1998) Consumption self-control by rationing purchase quantities of virtue and vice. *Mark Sci* 17:317–337.
- Ariely D, Wertenbroch K (2002) Procrastination, deadlines, and performance: Self-control by precommitment. *Psychol Sci* 13:219–224.
- Ashraf N, Karlan D, Yin W (2006) Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *Q J Econ* 121:673–697.
- Cialdini RB, Goldstein NJ (2004) Social influence: Compliance and conformity. *Annu Rev Psychol* 55:591–621.
- Cialdini RB, Trost MR (1998) Social influence: Social norms, conformity, and compliance. *The Handbook of Social Psychology*, eds Gilbert D, Fiske S, Lindzey G (McGraw-Hill, Boston), 4th Ed, pp 151–192.
- Lieberman N, Trope Y (2010) Construal-level theory of psychological distance. *Psychol Rev* 117:440–463.
- Agerström J, Björklund F (2009) Temporal distance and moral concerns: Future morally questionable behavior is perceived as more wrong and evokes stronger prosocial intentions. *Basic Appl Soc Psych* 31:49–59.
- Agerström J, Björklund F (2009) Moral concerns are greater for temporally distant events and are moderated by value strength. *Soc Cogn* 27:261–282.
- Gong H, Medin D (2012) Construal levels and moral judgment: Some complications. *Judgm Decis Mak* 7:628–638.
- Henderson MD, Trope Y, Carnevale PJ (2006) Negotiation from a near and distant time perspective. *J Pers Soc Psychol* 91:712–729.
- De Dreu C, Giacomanonio M, Shalvi S, Sligte D (2009) Getting stuck or stepping back: Effects of obstacles in the negotiation of creative solutions. *J Exp Soc Psychol* 45:542–548.
- Fujita K, Trope Y, Liberman N, Levin-Sagi M (2006) Construal levels and self-control. *J Pers Soc Psychol* 90:351–367.
- Giacomanonio M, De Dreu C, Shalvi S, Sligte D, Leder S (2010) Psychological distance boosts value-behavior correspondence in ultimatum bargaining and integrative negotiation. *J Exp Soc Psychol* 46:824–829.
- Van Lange P (1999) The pursuit of joint outcomes and equality in outcomes: An integrative model of social value orientations. *J Pers Soc Psychol* 77:377–349.
- Pruitt D, Kimmel M (1977) Twenty years of experimental gaming: Critique, synthesis, and suggestions for the future. *Annu Rev Psychol* 28:363–392.
- Hoffman E, McCabe K, Smith V (1996) Social distance and other-regarding behavior in dictator games. *Am Econ Rev* 86:653–660.
- Endsley MR (2017) From here to autonomy. *Hum Factors* 59:5–27.
- Blount S (1995) When social outcomes aren't fair: The effect of causal attributions on preferences. *Organ Behav Hum Decis Process* 63:131–144.
- Kircher T, et al. (2009) Online mentalising investigated with functional MRI. *Neurosci Lett* 454:176–181.
- Sanfey AG, Rilling JK, Aronson JA, Nystrom LE, Cohen JD (2003) The neural basis of economic decision-making in the ultimatum game. *Science* 300:1755–1758.
- Murphy R, Ackerman K, Handgraaf M (2011) Measuring social value orientation. *Judgm Decis Mak* 6:771–781.
- Paolacci G, Chandler J, Ipeirotis P (2010) Running experiments on Amazon Mechanical Turk. *Judgm Decis Mak* 5:411–419.
- Simpson B (2003) Sex, fear, and greed: A social dilemma analysis of gender and cooperation. *Soc Forces* 82:35–52.
- Van Vugt M, De Cremer D, Janssen DP (2007) Gender differences in cooperation and competition: The male-warrior hypothesis. *Psychol Sci* 18:19–23.
- Smith V (2003) Constructivist and ecological rationality in economics. *Am Econ Rev* 93: 465–508.
- Preacher KJ, Hayes AF (2008) Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behav Res Methods* 40: 879–891.
- Mell J, Lucas G, Gratch J (2018) Welcome to the real world: How agent strategy increases human willingness to deceive. *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems* (International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC).