



Climate change: Behavioral responses from extreme events and delayed damages[☆]

Riccardo Ghidoni^{a,*}, Giacomo Calzolari^{b,c}, Marco Casari^{b,d}

^a Tilburg University, Netherlands

^b Department of Economics, University of Bologna, Piazza Scaravilli 2, 40124 Bologna, Italy

^c CEPR, United Kingdom

^d IZA, Germany

ARTICLE INFO

Available online 3 November 2017

JEL classification:

C70

C90

D03

Q54

Keywords:

Social dilemma

Experiments

Greenhouse gas

Pollution

ABSTRACT

Understanding how to sustain cooperation in the climate change global dilemma is crucial to mitigate its harmful consequences. Damages from climate change typically occur after long delays and can take the form of more frequent realizations of extreme and random events. These features generate a decoupling between emissions and their damages, which we study through a laboratory experiment. We find that some decision-makers respond to global emissions, as expected, while others respond to realized damages also when emissions are observable. On balance, the presence of delayed/stochastic consequences did not impair cooperation. However, we observed a worrisome increasing trend of emissions when damages hit with delay.

© 2017 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Although scientists have convincingly established a causal link between greenhouse gas emissions and global climate change (IPCC, 2014), the way in which citizens perceive the issue may be simply through the experience of damages. News headlines are generally on the consequences of extreme events such as record temperatures, hurricanes or flooding that are outcomes of pollution and affect specific geographical areas. Another peculiar feature of climate change is the lag built into the earth system between the polluting actions and the system's reaction in terms of climate-related human impacts. Both these features imply a decoupling between polluting actions and their consequences. An usually unspoken argument among politicians and climate change experts is that it will likely take one or more major disasters to motivate citizens and nations to jump start mitigation efforts. Suffering environmental stress may be what can trigger citizens into action to stop climate change more than national plans contemplating changes in emissions. This conjecture motivates our behavioral study.

We focus on the ability to reach ambitious mitigation policies through voluntaristic actions when no binding treaty is in place, such as for example with the scheduling of periodic encounters after the Paris Agreement (Tollefson, 2016). More precisely, we design a climate change game as a N -person voluntary public bad game where decision-makers repeatedly interact under a long-run horizon (Dutta and Radner, 2004; Calzolari et al., 2016). Each decision-maker decides on a level of emissions, which brings individual benefits from production and consumption but generates a negative externality to everyone in terms of climate damages. Cooperation entails limiting the level of emissions. Through a laboratory experiment we vary how damages occur across treatments and study its influence on the ability to cooperate. The damage function is one of the fundamental elements for evaluating alternative policies to cope with climate change (Nordhaus, 2010) and has been the focus of a recent debate calling for a need to rethink the way damage functions are designed within Integrated Assessments Models (Wagner and Weitzman, 2015; Stern, 2015). Here we target two critical dimensions of damage functions – the *random* and *delayed* relation between polluting actions and their consequences – because they could both affect the behavioral ability of decision-makers to cooperate. All our specifications of damage vary the riskiness or timing but keep constant the overall level in terms of expected present value. We do so to make easier the empirical comparison across treatments. In a *Stochastic treatment* the damage takes the form of a random accident, whose probability increases in the level of global emissions. This treatment models the consequences of emissions in terms of extreme events, like flooding, droughts, or hurricanes. The aim is not

[☆] We are thankful to seminar participants at the University of Bologna and two anonymous referees. Casari acknowledges financial support through Italian Ministry of Education, University, and Research (FIRB-Futuro in Ricerca grant No. RBF084L83).

* Corresponding author at: Tilburg School of Economics and Management, Department of Economics, P.O. Box 90153, 5000 LE Tilburg, Netherlands.

E-mail addresses: r.ghidoni@uvt.nl (R. Ghidoni), giacomo.calzolari@unibo.it (G. Calzolari), marco.casari@unibo.it (M. Casari).

to capture a global catastrophe but instead low probability-high impact events that hit a country. We contrast this setting with a *Control treatment* where the damage from climate change occurs deterministically in proportion of global emissions. In a *Delay treatment* the damage is deterministic but hits decision-makers with a delay of two rounds – unlike the other two treatments where current damages depend on current emissions.

While some aspects of the field nicely map into our experiment, we made three major simplifications in order to facilitate participants' understanding of the task and to ease the empirical identification of the effects of the different treatments. First, we model climate damages as a flow externality that linearly increases in emissions, although a more accurate function would be a stock externality with possible non-linearities between emissions and damages (Burke et al., 2015; Dannenberg et al., 2015). A previous experiment showed a negative effect of pollution persistence on the empirical levels of cooperation in the long-run (Calzolari et al., 2016).¹ Second, we consider a limited number of players. Third, we include the deep income inequalities that exist in the field (Nordhaus, 2010; Tavoni et al., 2011) by having two types of participants, rich and poor, who simply differ in their private benefits from emissions.

In all our treatments, monitoring is perfect. After each round of play decision-makers can observe individual emission choices and damages of everyone else. These are propitious circumstances for cooperation to emerge. Under a long-run horizon – like the one considered here – the mitigation of damages may in fact realize under the threat of a punishment activated with the observation of an unexpected increase in others' emissions (the folk theorem, e.g. Fudenberg and Maskin, 1986). Such theoretical result would assume that all individuals follow strategies based on the observation of actions, i.e. emissions. However, individuals may in practice adopt strategies that react to experienced damages rather than actions. The reason may be behavioral, either related to salience or the cognitive costs to process information. On the one hand, damages directly influence payoffs and thus could be more salient to the decision-maker. On the other hand, even when observable, actions have to be interpreted in terms of motivating intentions, particularly when decision-makers form heterogeneous beliefs.

To sum up, greenhouse gas emissions generate delayed, random damages and hence actions (emissions) can be decoupled from their consequences (damages). What motivates this study is the possibility that some decision-makers rely more on experienced damages than actions, which calls for an empirical analysis of how different damage specifications could produce different outcomes in terms of mitigation.

The major result of our experiment concerns the strategies employed by participants in sustaining a cooperative mitigation. We show that participants react both to emissions and damages. In particular, some participants react to the emissions of others, as suggested by a canonical trigger strategy. Other participants, instead, react only to the extreme events or to the realized damages. A third group of participants respond to both emissions of others and individual damages. In Section 7 we conjecture on how the presence of these different types of individuals can relate to the differences in the overall cooperation levels we detect, in particular the withstanding levels of cooperation with stochastic and delayed damages and the increasing trend of emissions in the latter treatment.

The paper proceeds as follow. Section 2 places the contribution within the context of the literature about experiments on climate change and long-run cooperation. Section 3 presents the formal setup and experimental design. Section 4 puts forward some theoretical considerations about equilibrium predictions. Section 5 explains how the experiment was run. Section 6 describes the main results about aggregate emissions and strategies, while Section 7 discusses the results, some policy implications and concludes.

¹ Another dimension of the damage function that we do not consider here is its inter-generational feature (Sherstyuk et al., 2016).

2. Related literature

We contribute to two branches of the literature, one on climate change and another about sustaining long-run cooperation.

There exists a small but growing experimental literature on mitigation policies for climate change.² Some experiments model climate change as a problem of sustaining cooperation when facing an emission thresholds that may activate a catastrophe, while others, including the present one, model it with an incremental damage from pollution. Among the former category, the pioneering study is Milinski et al. (2008), who show that a higher probability of a catastrophe reduces emissions in the presence of a known tipping point. This result becomes weaker if the location of the tipping point is random, and more so in case of ambiguity (Barrett and Dannenberg, 2012, 2014; Dannenberg et al., 2015). Income inequality and the ability to communicate also affect the frequency of avoiding a catastrophe: Tavoni et al. (2011) show that groups that manage to reduce inequality during the play are the most cooperative, especially when communication is possible.

The experiments with a gradual impact of pollution on damages are relatively more recent. Sherstyuk et al. (2016) compares overlapping generations versus long-lived agents and reports that cooperation is harder to sustain for overlapping generations; Pevnitskaya and Ryvkin (2013) contrasts finite and indefinite horizons and find that participants learn to cooperate faster in the former setting, although they experience a last round drop; finally, Calzolari et al. (2016) study pollution persistence in a dynamic setting and show that it does not hamper cooperation per se but report a declining trend of cooperation for higher stocks of pollution. The novelty in our experimental design is to decouple actions and their consequences on damages, which in most studies are instead associated and indistinguishable. Our aim is to uncover the behavioral responses in a setting that replicates these key features present in the field.

The contribution of our paper to the vast literature about sustaining cooperation in repeated games rests on the distinction and observability of actions (emissions) and their consequences (damages). When the “shadow of the future looms sufficiently large”, cooperative outcomes can be obtained, possibly also the socially optimal outcome, with strategies punishing actions that deviate from a cooperative norm (Friedman, 1971; Dal Bó and Fréchette, 2017). Beginning with Green and Porter (1984), Abreu et al. (1990), Fudenberg et al. (1994), and Dutta (1995), the standard folk theorem has been extended to the case in which decision-makers do not perfectly observe others' actions, either because actions are observed with delay, as in our Delay treatment, or because observability only refers to an imperfect signal, such as the accident realization in our Stochastic treatment. Applying these results, we experimentally show that although the temptation to deviate from cooperation is generally stronger for strategies based on damages than emissions, cooperation could still be sustained when participants value sufficiently the payoffs from future interactions.

Some experimental papers on cooperation are related to our study. Bereby-Meyer and Roth (2006) study a repeated game with observable actions where outcomes can be either deterministic or probabilistic, depending on treatments. Relying on the psychological concept of “reinforcement” (Robbins, 1971), they report how a deterministic environment, granting a systematic reinforcement in the learning process, fosters cooperation as compared with the partial reinforcement available with random outcomes. Fudenberg et al. (2012) study the effects on cooperation of errors in implementing intended actions. They show considerable diversity in strategies, as we document in our analysis, and that successful strategies are “lenient” and “forgiving”: unexpected actions are not immediately punished, with attempts to restore cooperation. Camera and Casari (2009) manipulate monitoring of individual

² Although experiments on climate change face challenges of external validity, they play an important role integrating and complementing theory and field data. One definitive advantage of laboratory experiments is the possibility to control the environment and manipulate parameters, which enables to identify causal effects (Falk and Heckman, 2009).

histories and aggregate information on past cooperation that selectively add and remove the possibility to retaliate or adopt various punishment strategies. Finally, Nicklisch et al. (2016) experimentally find that when participants can jointly reduce the probability of a common stochastic damage, cooperation is enhanced. We confirm and extend this result to the extent that our stochastic damages are individual and participants have the possibility to observe emissions as well. In both cases participants appear to assess others' behavior with an ex post perspective, i.e. considering also the realization of outcomes.

3. Experimental design

We model climate change as a repeated social dilemma under three treatments – Control, Delay, and Stochastic – that vary the form taken by damages from the pollution externality. In a group of $N = 4$ decision-makers, everyone simultaneously takes a decision in every round $t = 1, 2, \dots$ over how much to emit, $e_t = (1, 2, \dots, 18)$. Individual payoffs are the difference between a benefit and a damage function:

$$\pi_i \equiv \text{Benefits}_i(e_i) - \frac{1}{N} \text{Damages}_i(E) \tag{1}$$

where $E = \sum_{j=1}^N e_j$ is the global emissions.

The benefit of an extra unit of emissions is private as it falls entirely on the decision-maker, while only $1/N$ of the damage does. Hence, emissions generate a negative externality on others in the group. There are four modifications with respect to the usual public good experiment, which make our framework similar to the model of Dutta and Radner (2004) as for payoffs.³ First, the game is framed as a public bad where the public project is fully provided by default and every unit of emission corresponds to moving contributions away from the “group account” into the “private account”. Second, the theoretical benchmarks of the one-shot Nash equilibrium and the socially optimal emission are not on the boundary of the action space, which is a desirable feature of an experimental design (Laury and Holt, 2008). Our benefit function is non-linear in emissions, as additional units have a lower return, while the damage function is linear. As we will see, this generates an interior Nash at 12, which is far from the upper bound of 18 and allows for anti-social behavior. Moreover, the socially optimal level of emission is at 3. Third, to mimic GDP inequality in the world arena, we introduced payoff heterogeneity within the group, with rich decision-makers enjoying a higher return from the private account (i.e. the benefit function) than poor ones while suffering identical levels of damages. More precisely, the benefit function is, for a level of emission $e(t)$ at time t :

$$\text{Benefits}_i(t) \equiv 100 \ln(a_i e_i(t)) \tag{2}$$

The parameter a_i is set at 40.05 for half of the group members (rich) and 8.01 for the others (poor). This asymmetry in a_i could capture technological differences in carbon intensity leading decision-makers to achieve different benefits for the same level of emission.⁴ Fourth,

³ Dutta and Radner (2004) study dynamically persistent emissions that accumulate in a stock over time, while in our analysis there is no persistence (Calzolari et al., 2016, experimentally study pollution persistence).

⁴ Both types of decision-makers have the same emission capacity. To ensure rich and poor decision-makers have the same social optimum and stage-game Nash equilibrium (see Section 4) and ease empirical comparisons, the gap between rich and poor decision-makers is modeled as a gap in private benefits (Eq. (2)). While this is a strong simplification, the experiment roughly reflects stylized facts from IPCC (2014) and the RICE model (Nordhaus, 2010). Rich decision-makers mirror high income countries with a per capita GNI above \$12,745 (World Bank threshold in 2010), whose GHG emissions amounted to 18.7Gt in 2010 (IPCC, 2014). Instead, poor decision-makers approximately resemble countries with a per capita GNI lower than \$12,745: upper-middle income countries' emissions were quite close to high income countries' emissions (18.3Gt), emissions from low and lower-middle income countries were instead lower (11.3Gt). When focusing on the regions of the RICE model, rich regions have an average GNI per capita 4.8 times higher than poor regions. Poor regions are Africa, China, Eurasia, India, and Other Asia ($N = 5$, average GNI per capita = \$7125.9); rich regions are EU, Japan, Latin America, Middle East, Russia, USA, and Other High Income ($N = 7$, average GNI per capita = \$34,085).

we implement a long-run horizon to capture the long life of state entities and of the climate change problem. In the lab, the interaction is indefinite and is implemented through a random stopping rule. After every round there is a random draw: an additional round is played with probability $\delta = 0.92$ and the sequence stops with probability $(1 - \delta)$. As a consequence, the length of a sequence is variable and nobody knows when the last round will take place. The “shadow of the future” remains the same as the rounds proceed because the continuation probability δ is constant and common knowledge. Such probability can be interpreted as the discount factor of a risk-neutral decision-maker who lives forever.

While the benefit function is identical in all treatments, the damage function is treatment-specific (Table 1). In the *Control treatment* damages from global emissions are deterministic and hit immediately in the same round of emissions according to the following damage function:

$$\text{Damages}_i(t) \equiv c1 \times E(t) \tag{3}$$

where the parameter $c1 = 33.375$ determines the magnitude of the damage for each unit of emissions. Damages are proportional to emissions to keep the design simple.⁵

In the *Delay treatment*, the damages are also deterministic but hit with a delay of two rounds. As a consequence, there will be no damages in the first two rounds:

$$\text{Damages}_i(t) \equiv \begin{cases} 0, & \text{if } t = 1, 2 \\ c2 \times E(t-2), & \text{if } t > 2 \end{cases} \tag{4}$$

The damage parameter $c2 = 39.432$ is set taking as reference the value in the Control treatment, $c2 = \delta^2 c1$, so to keep the same present value for the damage generated by one unit of emissions.⁶

Finally, in the *Stochastic treatment*, damages hit immediately but at random: a fixed accident of magnitude $K = 830$ may hit one or more decision-makers with a probability which linearly increases at a constant rate of about 1 percentage point for every unit of global emissions ($\alpha = 0.01005$):

$$\text{Probability of an accident}_i(t) \equiv \alpha \times E(t) \tag{5}$$

The accident's probability ranges from a minimum of 0.0402 if everyone emits 1 through 0.7236 if everyone emits 18.⁷ By design there is no way to reduce accident's risk to zero and, no matter how high emissions are, the accident always remains uncertain. All group members share an identical risk of suffering an accident as the probability depends on the global rather than individual emissions. However, there are independent draws for each decision-maker to determine if an accident occurs. Hence, the damage level will be identical across group members only in event of zero or N accidents and will differ in all other random events. In expectation, the marginal damage from a unit of emissions is similar to the Control treatment, $\alpha \times N \times K = c1$.⁸

There are many alternative ways to incorporate the randomness of climate change into the design. Through the Stochastic treatment we aim to model extreme events rather than global catastrophes.

⁵ As already mentioned, in the field damages are likely to be a convex function of temperatures (Burke et al., 2015) but in theoretical models others have also employed a linear approximation (Dutta and Radner, 2004). Moreover, pollution persistence is not included in the model to simplify the design.

⁶ This calibration leaves the stage-game Nash and the socially optimal levels of emissions unaffected under the assumption of risk-neutral decision-makers.

⁷ We modeled extreme events through a linear function rather than a Pareto distribution to make it simple for participants to understand the environment. Adopting cooperative strategies allows participants to induce a fairly low probability of an accident.

⁸ The equivalence is almost exact as the value is 33,366 in Stochastic and 33,375 in Control. This calibration leaves the Nash and socially optimal levels of emissions unaffected under the assumption of risk-neutral decision-makers.

Table 1
Overview of the experiment.

Parameters		Control	Delay	Stochastic
$D(\sum e_i)$	Damage function	Deterministic and immediate ($c1 = 33.375$)	Deterministic and delay of two rounds ($c2 = 39.432$)	Immediate random accident of $K = 830$ ($\alpha = 0.01005$)
δ	Discount factor (continuation probability)	0.92	0.92	0.92
ar	Benefit parameter for rich decision-maker	40.05	40.05	40.05
ap	Benefit parameter for poor decision-maker	8.01	8.01	8.01
Predictions and results				
e^{**}	Social optimum (individual emission)	3	3	3
e^*	Nash equilibrium	12	12	12
e_i	Result: average emission (range 1–18)	9.4	7.9	8.5
Sessions (dd/mm/yy)				
		20/05/15	16/03/16	23/06/16
		21/05/15	17/03/16	24/06/16
		27/05/15	18/03/16	27/06/16
Number of participants (main task + side task)				
		60 + 15	60 + 15	60 + 13
Number of groups				
		55	45	45
Number of sequences				
		11	9	9
Average length of a sequence				
		10.8	16.8	13.8

Notes: Data from sessions in the Control treatment have also been analyzed in a related paper (Calzolari et al., 2016, Immediate treatment). Average emissions are computed as the mean of the individual emissions in a group in a sequence. For time constraints, sessions 20/5/2015 (Control), 17/03/2016 (Delay), and 24/06/2016 (Stochastic) were interrupted during the third sequence following the protocol described in footnote 14; session 18/03/2016 (Delay) was interrupted during the second sequence. In session 24/06/2016 (Stochastic), only 23 volunteers showed up, so two participants to the side task were missing.

While a global catastrophe causes similarly losses to all players, extreme events such as hurricanes tend to hit areas asymmetrically. This original feature sets this study apart from the previous climate change experiments. Furthermore, it facilitates a cleaner empirical identification of individual-level effects: a common shock to all participants would limit the variation of impacts and hence restrict the possibility to identify individual strategies, which is a main goal.

In the Stochastic treatment, before the climate game, we elicited the risk preferences of all participants following the design of Karle et al. (2015). In particular, participants were administered two tasks, one in the gain domain and the other in the loss domain. In the former task, participants had to make six binary choices. Each decision was between a 50–50 lottery yielding either 0 or 3€, and a certain amount (0.3, 0.6, 0.9, 1.2, 1.5, or 1.8€). The latter task was similar except for that the certain amount was always 0€ and the lottery either paid 3€ or involved a loss (−0.3, −0.6, −0.9, −1.5, −2.1, or −3€). One of these twelve decisions was randomly drawn at the end of the session, and participants were paid accordingly. Participants did not receive any feedback on the lotteries outcomes until the end of the session.

In all treatments, the present expected value of a decision-maker i 's current and future payoffs is,

$$\Pi_i = \sum_{t=0}^{\infty} \delta^t \pi_i(t). \quad (6)$$

4. Theoretical benchmarks

This Section provides the theoretical benchmarks which will be useful to evaluate and interpret the experimental results. We proceed in three steps. In step one we identify the socially optimal level of emissions and in step two we present the level of emissions in the one-shot equilibrium with decentralized choices. The contrast between the two levels of emissions highlights the social dilemma dimension of the climate game. In step three we characterize some relevant equilibria of the repeated game. According to the standard folk theorem (Friedman, 1971), when the shadow of the future looms sufficiently large and monitoring is perfect, decision-makers can adopt strategies that support cooperative outcomes, possibly also the socially optimal one. These strategies can take the form of grim triggers where decision-makers contemplate permanent punishment when they observe a deviation from a cooperative norm. The punishment is collective because in our setting it is impossible to target a single

decision-maker. As usual, a multiplicity of equilibria arise, hence coordination is a relevant empirical issue. We assume that all decision-makers are risk neutral.

If decision-makers cooperate, they maximize the unweighted sum of individual present-valued payoffs,

$$\Pi = \frac{N}{2} (\Pi_r + \Pi_p).$$

They set a time-invariant **socially optimal emission** $e^{**} = 3$, where the marginal benefit from the individual emission, $100/e_i$, equals to the marginal damage caused on the whole group. In the Control treatment the marginal group's damage is $N \times \frac{c1}{N}$. In the other two treatments – given our parametrization – the marginal group's damage is equal, in expectation, to the level in the Control treatment and hence the socially optimal emission is also at $e^{**} = 3$.

When decision-makers act independently there always exists an equilibrium in which the level of emissions in any round corresponds to the Nash equilibrium of the one-shot stage-game. Here, each decision-maker equates the marginal benefit from her individual emission to the individual marginal damage, which is a fraction $\frac{1}{N}$ of the group's damage. As in standard public goods games, this condition does not depend on others' emissions, hence the **one-shot Nash equilibrium emission** $e^* = 12$ is unique, and is equal to $N \times e^{**}$ in all treatments.

We now turn to other levels of emission that can be supported in the repeated game, distinguishing between strategies based on the observation of others' emissions and strategies based on damages suffered.⁹

4.1. "Observational equilibria"

Here we consider equilibria supported by strategies that are based on the past individual emissions of all N decision-makers as observed at the end of each round, which are the most common class of strategies in the folk theorem literature. Observability allows decision-makers to use trigger strategies that contemplate a punishment upon observing levels of emission that are interpreted as deviations. One can easily prove the following.

⁹ In Section 7 we will discuss the consequences of the interaction between decision-makers who base their strategies on observed emissions and others on damages.

Remark 1. In all treatments, if decision-makers are “observational”, they can support in equilibrium any level of individual emission between $e^{**} = 3$ and $e^* = 12$.

A proof of Remark 1 hinges on the canonical grim trigger strategies. Consider the Control treatment first. Suppose $N - 1$ decision-makers rely on a strategy that contemplates emitting a low level $3 \leq \hat{e} < 12$ if this is what happened in the past rounds and, instead, a permanent reversion to the Nash equilibrium emission e^* if they observe an individual emission different from \hat{e} . One can show that at any round t , a decision-maker prefers to keep low emissions \hat{e} instead of (the optimal deviation) e^* . The present value payoffs of emitting \hat{e} is,

$$\Pi_i = \frac{1}{1-\delta} \left[100 \ln(a_i \hat{e}) - \frac{c1}{N} \hat{e} \times N \right]. \tag{7}$$

Alternatively, the payoff for emission e^* is the sum of the current round payoff when everyone else emits \hat{e} and the future rounds payoffs when everyone else enters the punishment mode and also emits e^* ,

$$\begin{aligned} \Pi_i = & \left[100 \ln(a_i e^*) - \frac{c1}{N} (3(N-1) \times \hat{e} + e^*) \right] \\ & + \frac{\delta}{1-\delta} \left[100 \ln(a_i e^*) - \frac{c1}{N} e^* \times N \right] \end{aligned} \tag{8}$$

The payoff Π_i in Expression (7) is always larger than that in Expression (8) if the decision-maker is sufficiently patient. With the parameter values set in the experiment, the condition that guarantees this preference if we want to support the socially optimal outcome $\hat{e} = e^{**}$ is a discount factor δ above the critical threshold $\bar{\delta} \approx 0.3$. Such condition is well satisfied in the experiment given that $\delta = 0.92$.¹⁰

The same reasoning applies to the Stochastic treatment because it is isomorph to the Control treatment. The proof for the Delay treatment is also similar, except that the cost of punishment hits the deviator only after three rounds. In fact, in the round of the deviation the other decision-makers are taken by surprise and start punishing the following round, with consequences accruing after two additional rounds. Although here the punishment is less effective because it hits with delay, this treatment admits the socially optimal outcome as an equilibrium for a discount factor higher than $\bar{\delta} \approx 0.7$, which is larger than in the other treatments but still smaller than $\delta = 0.92$.

4.2. “Experiential equilibria”

We now consider all decision-makers who follow strategies exclusively based on realized damages rather than global emissions. In the Control treatment, the distinction between observational and experiential equilibria is immaterial since there is no decoupling between actions and damages.

In the Stochastic treatment, the experiential strategy is based on the realized accidents. We first consider the case of decision-makers keeping track of all realized accidents in their group, and then we briefly move to the case of decision-makers keeping track only of their own accidents.¹¹

Unlike with “observational” decision-maker, “experiential” decision-makers can never be sure that a deviation has effectively occurred in the group. We define $A(t)$ as the event in which at least one accident occurred in round t and $Pr(A|\hat{E})$ as its probability for a given level of global

emissions $\hat{E} = \hat{e} \times N$. A cooperative outcome can be sustained with a punishment mode that lasts for a finite number T of rounds (instead of being permanent). When no accidents have occurred, decision-makers emit \hat{e} . Instead, when at least one accident has occurred in the group, they temporarily emit e^* for the next T rounds, regardless of additional accidents; this is the so-called “quasi-punishment” phase.¹² On the equilibrium path the expected payoff is,

$$\begin{aligned} \pi_i = & 100 \ln(a_i \hat{e}) - \alpha \hat{E} K + \delta \left(1 - Pr(A|\hat{E}) \right) \pi_i \\ & + \delta Pr(A|\hat{E}) \left\{ \frac{1-d^T}{1-\delta} [100 \ln(a_i e^*) - \alpha N e^* K] + \delta^{T+1} \pi_i \right\}. \end{aligned} \tag{9}$$

Here, we construct the experiential equilibria as if monitoring was imperfect due to unobserved emission, although in the experiment they were observable. Under imperfect monitoring, decision-makers do not know for sure if the realization of an accident is the consequence of a deviation or not and this will trigger some high emissions e^* also along the equilibrium path (the term in curly brackets). Moreover, a deviation may go “unnoticed” unless it triggers an accident, in which case it induces a continuation payoff that is the same as the one along the equilibrium path (the same curly brackets just described).

Notwithstanding this reduced incentive power of punishments, one can show that there exists a (decreasing) function $T(\hat{e})$, such that π_i is larger than the payoff associated with a deviation if $T \geq T(\hat{e})$. By keeping emissions low at $\hat{e} < e^*$, decision makers keep the probability to trigger the quasi-punishment low. The longer the punishment phase, the more efficient is the emission level that can be implemented ($\hat{e} = e^*$ for $T \geq 10$).

Hence, the type of cooperation reached by experiential decision-makers contemplates higher emissions in some rounds triggered by the realization of stochastic accidents. Even if in equilibrium decision-makers are able to sustain a level of individual emission \hat{e} without accidents, they will switch to T rounds of quasi-punishment with higher emissions after the realization of any accident and end up with an average level of emissions well above \hat{e} . This implies that, even in the most cooperative scenario, experiential decision makers will emit on average more than the socially optimal emission e^* .

When decision-makers consider their own accidents only and disregard those of the others, what matters is the probability $Pr(A_i|\hat{E})$ that decision-maker i experiences an accident. The enhanced difficulty in sustaining cooperation is that quasi-punishment rounds are here asynchronous because different decision-makers care for different and independent accidents. A quasi-punishment phase may thus trigger accidents to other decision-makers who, in turn, activate their own quasi-punishments propagating even higher emissions. Although an explicit derivation of an equilibrium would considerably complicate the analysis, one can see that the difficulty to jointly identify and react to deviations makes cooperation weaker although not impossible. What is relevant to us is that in any case decision-makers would individually react by increasing emissions when an individual accident has occurred (Yamamoto, 2012).

Finally, consider the Delay treatment. Here, experiential decision-makers can realize that a deviation occurred by simply inspecting the current damage. However, since this observed deviation refers to two preceding rounds, their punishment begins with a delay with respect to observational decision-makers. More precisely, a deviation is detected two rounds later so that in the second, third and fourth rounds after deviation the other decision-makers’ emission are still lower at \hat{e} . Only from the fifth round onward the deviator is hit by the punishment and all decision-makers revert to the Nash stage-game emissions e^* . The

¹⁰ Emission levels larger than 12 should not occur in equilibrium because individually and collectively dominated by e^* .

¹¹ Recall that in our experiment others’ accidents are observable, and so are individual emissions. If decision-makers keep track of all individual realized damages, the game is one with “imperfect public information” (Fudenberg et al., 1994). If instead decision-makers disregard the realizations of others’ accidents, then the Stochastic treatment becomes a complex game of “imperfect private monitoring” where the possibility to obtain cooperation via a folk theorem argument is limited (Yamamoto, 2012).

¹² With our parametrization, the probability of an accident is bounded away from zero and from one for any level of global emission. This implies that temporary punishment upon the observation of a damage is necessary to support more efficient outcomes in our game with “full support”, as in Green and Porter (1984).




Fig. 1. Average emission by treatment. *Notes:* The unit of observation is a group in a sequence ($N = 55$ in Control, $N = 45$ in Delay, $N = 45$ in Stochastic). Individual emissions can range from 1 through 18. The vertical segments represent the 95% confidence interval. The red-upper and the green-lower horizontal lines respectively indicate the Nash individual emission of the stage-game ($e^* = 12$) and the socially optimal level of individual emissions ($e^s = 3$).

logic for the possibility to support emissions more cooperative than e^* is the same as with no delayed damages, except for the “diluted” efficacy of the punishment. Observational decision-makers can support the socially optimal outcome in the Delay treatment if sufficiently patient, with a threshold value for δ now being $\bar{\delta} = 0.84$.

We can now summarize the following theoretical results.

Remark 2. *Experiential decision-makers react to damages and, although they are slower to react to deviations than observational decision-makers, they may still be able to cooperate reducing emissions. (i) In the Stochastic treatment, they increase emissions after realized individual or collective accidents. (ii) In the Delay treatment, they increase emissions reacting to damages of two-rounds previous emissions.*

5. Experimental procedures

We have run 9 sessions at the University of Bologna, with a total of 180 participants. Procedures aimed at ensuring that all participants had a good level of understanding of the instructions (see Appendix B for an English translation). To this end, in every session we recruited 25 participants but only 20 were actually performing the main task: the selection was based on a quiz about the instruction.¹³ There was a sequence of “dry runs” played against robots that were varying their emission level round after round. A session comprised three or four sequences of interaction with monetary incentives.¹⁴ After

¹³ The excluded participants had to do a side task with a flat payment of 0.50€ per round plus a show-up fee of 5€.

¹⁴ Participants were recruited for up to three hours and a half. For long sessions (more than two hours and forty minutes), we informed participants that the current sequence of interaction was the last one and that the experiment would end within thirty minutes. In this case the exact termination moment was random, as we explained to the participants, with a random draw between 1 and 30. In one session of the Delay treatment, the session was terminated during the second sequence due to time constraints. Since long sessions were randomly and unexpectedly interrupted, this should have no impact on participants' behavior. We did not conduct any ex post debriefing to limit the duration of the session.

every sequence, all participants were rematched with completely different people to play the next sequence (perfect stranger matching protocol).

Table 2
Treatment effects on the average emission.

Dependent variable:	Control vs. Delay		Control vs. Stochastic	
	(1)	(2)	(3)	(4)
Average emission in a group	First round	All rounds	First round	All rounds
Treatment dummies				
Delay	-1.289** (0.541)	-1.213* (0.661)		
Stochastic			0.252 (0.653)	-0.560 (0.597)
Sequence dummies				
Sequence 2	0.748 (0.652)	0.022 (0.745)	1.085 (0.742)	1.261* (0.699)
Sequence 3	1.022 (0.688)	-0.072 (0.787)	1.338* (0.784)	1.125 (0.716)
Sequence 4	-0.036 (0.847)	0.457 (0.965)		
Length of past sequence	-0.147*** (0.046)	-0.135** (0.057)	-0.097** (0.042)	-0.068* (0.038)
Length of current sequence		0.060*** (0.023)		0.035 (0.022)
Constant	8.968*** (0.709)	9.727*** (0.900)	8.248*** (0.722)	8.721*** (0.724)
Observations	100	100	90	90

Notes: Results from OLS regressions are reported. The unit of observation is a group in a sequence. Variables “Delay” and “Stochastic” are dummies respectively taking value 1 in the Delay and Stochastic treatments, and 0 in the Control treatment. The variable “Length of past sequence” counts the number of rounds in the previous sequence; in sequence 1 it is set to 12.5.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

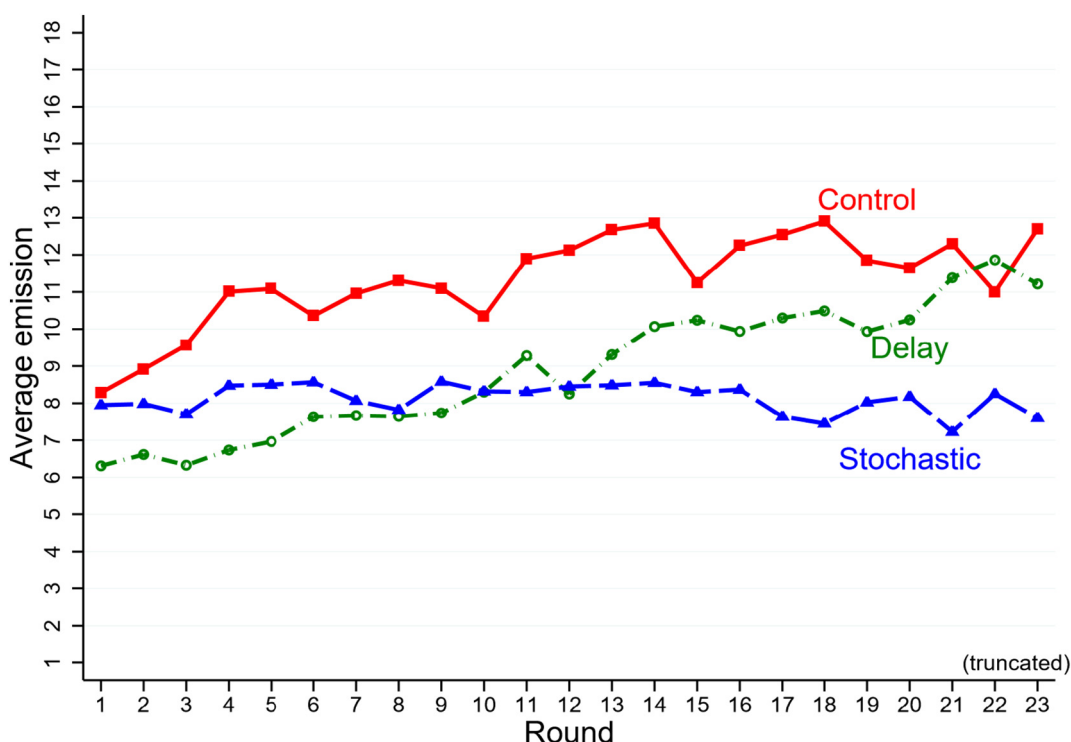


Fig. 2. Average emission over rounds. *Notes:* The unit of observation is a group in a sequence. In round 1 the number of observations is $N = 55$ in Control, $N = 45$ in Delay, $N = 45$ in Stochastic; in round 23 the number of observation is $N = 5$ in Control, $N = 15$ in Delay, $N = 10$ in Stochastic. Individual emissions can range from 1 through 18 with the socially optimal level $e^* = 3$ and the Nash equilibrium of the stage-game $e^* = 12$.

6. Results

We report six main results, some about aggregate outcomes (Results 1–2) and others about the strategies followed by participants (Results 3–6).

6.1. Aggregate results

Result 1. (Aggregate cooperation). *Delayed damages lower aggregate emissions and Stochastic damages do it to a marginal extent.*

Support for Result 1 comes from Fig. 1 and Table 2. Fig. 1 shows that the average emission is 7.9 in Delay, which is statistically significantly < 9.4 in Control both according to a non-parametric test (Wilcoxon-Mann-Whitney test: p -value = 0.011, $N_C = 55$, $N_D = 45$) and OLS regressions (Table 2, col. 1 and 2). The evidence for the Stochastic treatment is somewhat weaker in terms of magnitude (8.5 vs. 9.4). Differences in emissions are statistically significant between Control and Stochastic according to a non-parametric test (Wilcoxon-Mann-Whitney test: p -value = 0.036, $N_C = 55$, $N_S = 45$) but not to OLS regressions (Table 2, col. 3 and 4). The unit of observation in the regressions of Table 2 is a group in a sequence and we control for sequences order and length. After checking for heterogeneous responses, we will discuss these observations in the concluding Section.

Result 2. (Time trends). *With delayed damages, emissions exhibit a steadily increasing trend over the rounds. No clear trend emerges in Control and Stochastic treatments.*

Support for Result 2 comes from Fig. 2 and Table 3. Fig. 2 illustrates the emissions trend within a sequence. The Delay treatment starts with emissions that are remarkably lower than in Control (Fig. 1 and Table 2, col. 1) and then emissions steadily increase over the rounds. This trend is confirmed by an OLS regression explaining individual emission choices (Table 3, col. 2) that controls for a host of factors

such as sequence order and length, rich vs. poor type, level of understanding of the instructions, and limited liability issues.¹⁵ For the Control treatment, Fig. 2 shows an upward tendency that is not statistically significant (Table 3, col. 1).¹⁶ No trend emerges in the Stochastic treatment (Fig. 2 and Table 3, col. 3).

Two noteworthy patterns emerge from the data about inequality and risk preferences. On average, poor participants emit more than rich ones in every treatment (9.7 vs. 9.1 in Control, 8.1 vs. 7.6 in Delay, and 9.1 vs. 7.9 in Stochastic), but the difference in aggregate behavior is statistically significant only in the Delay and Stochastic treatments according to non-parametric tests (two-sided sign tests: Control: p -value = 0.892, $N_R = N_P = 55$, Delay: p -value = 0.073, $N_R = N_P = 45$, Stochastic: p -value = 0.036, $N_R = N_P = 45$). The evidence is more mixed when using OLS regressions (Table 3, rich participant dummy).¹⁷

For the Stochastic treatment we can also evaluate the role of individual risk preferences over emission choices. When risk attitude is elicited in the gain domain (the precise definition of the dummies is in the note of Table 3), it is not significantly correlated with emissions (Table 3, col. 3). Instead, when risk attitude is elicited in the loss domain, there is a statistically significant relation with emissions: risk seeking participants emit on average more than those that are neutral.

¹⁵ In some observations a participant ended up with negative cumulate earnings. In this case limited liability may have played a role for their subsequent actions. These observations were 9.9% in Control, 7% in Delay, and 6% in Stochastic. The unit of observation is a participant's choice in a round. Limited liability occurs if a participants' show-up fee and cumulate earning over the session ends up below 10€.

¹⁶ To reconcile the apparent differences in average emissions reported in Figs. 1 and 2 recall that the indefinite horizon naturally generates a declining number of observations (e.g. in our Control treatment, in round 23 there are five groups only). Therefore, observations in the last rounds "weight" much less than those in the first rounds when calculating overall average emissions.

¹⁷ Despite this, the highest level of inequality is reached in Delay (Gini coefficient = 0.383) compared to Control (0.529) and Stochastic (0.466). The unit of observation is one participant's total earnings at the end of a sequence.

Table 3
Regressions of individual emission.

Dependent variable:	(1)	(2)	(3)
Individual emission in the current round	Control	Delay	Stochastic
Time trend within a sequence (round)	−0.005 (0.031)	0.092** (0.038)	0.016 (0.028)
Sequence dummies			
Sequence 2	−2.447 (1.602)	−0.592 (1.145)	1.663** (0.690)
Sequence 3	−0.208 (1.372)	−1.028 (1.067)	0.028 (0.771)
Sequence 4	−1.001 (1.474)	1.680 (1.628)	
Length of past sequence	−0.137 (0.132)	−0.082 (0.086)	−0.115*** (0.033)
Rich participant dummy	0.919 (0.643)	−0.268 (0.553)	−0.943 (0.761)
Mistakes in the quiz	0.176 (0.325)	0.280 (0.256)	0.775** (0.352)
Limited liability	3.472*** (1.227)	5.086** (2.214)	1.737*** (0.558)
Risk averse in the gain domain			0.032 (0.424)
Risk seeking in the gain domain			0.688 (0.605)
Risk averse in the loss domain			−0.466 (0.411)
Risk seeking in the loss domain			1.053* (0.541)
Constant	11.223*** (1.803)	7.916*** (1.249)	8.121*** (0.791)
Observations	2380	3020	2480
R2	0.0671	0.2062	0.0746

Notes: Results from OLS regressions are reported. The unit of observation is a participant's emission choice in a round. Standard errors are clustered at the level of a group in a sequence. The variable "Length of past sequence" counts the number of rounds in the previous sequence; in sequence 1 it is set to 12.5. The variable "Mistakes in the quiz" counts the number of mistakes that a participant made in the quiz on the instructions. The variable "Limited liability" is a dummy taking value 1 if the emission decision was made under limited liability, and 0 otherwise. The dummy "Risk averse in the gain domain" is equal to 1 if the participant chose the lottery against the certain positive amount less than three times. The dummy "Risk seeking in the gain domain" is equal to 1 if the participant chose the lottery against the certain positive amount more than three times. Dummies "Risk averse in the loss domain" and "Risk seeking in the loss domain" are similarly defined. All risk dummies neglect whether the participant violated single crossing.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

6.2. Strategies of the representative participant

Our experimental design allows to go beyond the aggregate results about cooperation and to shed light on the type of strategies followed by participants in the repeated game. Here we study the strategies of the representative participants (Results 3–4) and in Section 6.3 we provide a simple classification of the individuals to further corroborate and specify the findings. The main theme of analysis is how a participant who may want to cooperate in reducing emissions reacted to a perceived defection. We begin with the study of observational strategies.

Result 3. (Observational strategies). *In the Control treatment, the representative participant responds to a perceived defection with a temporary increase in emissions.*

Support for Result 3 comes from Fig. 3 and Table 4. Data from the Control treatment suggest that when the representative participant observed high emissions by others in the group, she switched from a cooperative to a punishment mode. As seen in Section 4, an appropriately defined trigger strategy can sustain a fully cooperative equilibrium in our setting. While previous experiments with two players and two moves have already documented a similar pattern (Camera and

Table 4
Strategies of the representative participant in the Control treatment.

Dependent variable:	(1)	(2)	(3)	(4)
Individual emission in current round				
Trigger: average emission of others > 12				
Lag 1	1.711*** (0.443)	1.782*** (0.438)		1.759*** (0.413)
Lag 2	0.246 (0.336)	0.215 (0.339)		0.217 (0.312)
Lag 3		−0.004 (0.220)		0.213 (0.162)
Lag 4		0.206 (0.233)		0.310 (0.286)
Any previous round	0.407 (0.528)	0.381 (0.520)		0.863 (0.532)
Trigger: personal loss in round payoff				
Lag 1			1.147* (0.625)	0.029 (0.484)
Lag 2			0.592 (0.452)	0.302 (0.404)
Lag 3			−0.569 (0.439)	−0.735 (0.446)
Lag 4			−0.038 (0.289)	−0.196 (0.338)
Any previous round			−0.668 (0.637)	−1.447** (0.607)
Constant	3.968** (1.850)	3.996** (1.846)	4.058* (2.112)	4.836** (1.845)
Round dummies	Yes	Yes	Yes	Yes
Sequence dummies	Yes	Yes	Yes	Yes
Participant dummies	Yes	Yes	Yes	Yes
Length past sequence	Yes	Yes	Yes	Yes
Limited liability	Yes	Yes	Yes	Yes
Observations	2160	2160	2160	2160
R2	0.4453	0.4454	0.4248	0.4505
Adjusted R2	0.4126	0.4122	0.3903	0.4162

Notes: Results from OLS regressions are reported. The unit of observation is a participant emission decision in a round. Decisions in round 1 are dropped. Standard errors are clustered at group level.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

Casari, 2009; Dal Bó and Fréchette, 2017), the novelty of our result mainly lies in showing it in a N -person game with a multi-level action space.

Recall that, following a defection of some opponent, a trigger strategy involves a shift to a punishment mode with higher emissions for some number of subsequent rounds. For the Control treatment, the finding emerges from an OLS regression that explains individual emission choices using regressors that trace the strategy and a set of controls (Table 4, cols. 1 and 2). Controls include dummies for round, sequence, participant, and limited liability, as well as the length of the past sequence. In the regression model, we assume that a defection occurs if the emissions of the other three group members are on average equal or above 12, but we have checked other levels (see Table A.13 in Appendix). This analysis sheds light on the type of strategies employed by the representative participant generalizes that of Camera and Casari (2009) to N players and a multi-level action space. Although the way to code regressors in order to trace strategies is subject to some discretion, the approach has the advantage to detect whether participants followed theoretically well-known strategies, such as grim trigger or tit-for-tat.

The regressors that code the strategy aim to trace the response of the representative participant in the rounds that follow a perceived defection. We mostly focus to the four rounds after a defection by including four "Lag" regressors, which have a value of 1 only in one round following a defection and 0 otherwise. For example, the "Lag 1" regressor takes value 1 only in the round after the defection (0 otherwise). The "Lag 2" regressor takes value 1 only in the second round following a defection (0 otherwise). Similarly for the "Lag 3" and "Lag 4" regressors. However, we also consider a "grim trigger"

regressor labeled “Any previous round”, which has a value of 1 in all rounds following a defection and 0 otherwise.

Fig. 3 panel (a) illustrates the estimated reaction in emissions over the rounds, following an observed defection in the Control treatment. The illustration for round lags 1 through 4 is based on the sum of the coefficients of the grim trigger regressor and the lag regressor with the appropriate lag. The illustration for lag 5 is based on the effect of the grim trigger regressor only.¹⁸

The pulse pattern of response to an observed defection suggests a temporary downward shift in cooperation levels immediately after a defection. The lag 1 regressor is significantly different from zero, while the estimated coefficients of all other strategy regressors, including the grim trigger one, are not significantly different from zero (Table 4).

A pattern along the lines of Result 3 emerges also from the analyses of observational strategies in the other treatments. The finding comes from a similar estimation procedure carried out for the Delay and Stochastic treatments using the same emission threshold of 12.¹⁹

In the Delay treatment, the representative participant immediately increases emissions after an observed defection in a statistically significant way (lag 1, Table 5, col. 1); also in the Stochastic treatment there is a statistically significant immediate response (Table 6, col. 1). The main differences between Control and the other treatments seem to be (i) the presence of a more permanent punishment to a defection, as estimated by the “Any previous round” regressor (Tables 5 and 6, col. 1); (ii) a moderated pulse response to defections in the Delay and Stochastic treatments as compared with the Control treatment. Fig. 3 illustrates these pulse patterns of responses to an observed defection with the solid lines in panels (b) and (c) labeled “Observational”.

Result 4. (Experiential strategies). *The strategy of the representative participant responds both to the observed actions as well as to the experienced damage.*

Support for Result 4 comes from Fig. 3 and Tables 5–6. These findings emerge from the Stochastic and Delay treatments, where one can possibly decouple these observational and experiential strategies. Let’s begin with the evidence from the Stochastic treatment, where the empirical distinction between the two classes of strategies is more intuitive. We exploit the presence of random accidents, which determined a large shock on the current earnings, and tracked the reaction to them in terms of emissions of the representative participant. The empirical frequency of accidents in a group was as follow: in 2% of cases everyone in the group experienced an accident and in 21% of cases nobody experienced an accident. The mode was of one accident in the group in the round (38%). The data suggests that the representative participant increased emissions immediately after experiencing an accident (Table 6, col. 2). The reaction was statistically significant but temporary, i.e. limited to lag 1. As already mentioned, the estimate of observational strategies shows a strong immediate reaction (lag 1) and a smaller but permanent effect (Any previous round, Table 6, col. 1). We also performed a joint estimate of observational and experiential strategies and the patterns do not change substantially, with coefficients slightly smaller in magnitude (Table 6, col. 3). The two classes of strategies are illustrated with the two (solid and dashed) lines in Fig. 3 panel (c). Hence, the representative participant is responding with higher emissions both to others’ actions, when higher than a threshold, and also to personal payoffs shocks.

¹⁸ If at least one of the five strategy regressors estimated in Table 4 has a positive coefficient, then this could be the consequence of a representative participant switching from a cooperative to a punishment mode. We can illustrate this by the following example: a representative participant who punishes for exactly three rounds following a perceived defection generates estimated positive coefficients for the Lag 1, Lag 2, and Lag 3 regressors.

¹⁹ Robustness checks with alternative thresholds can be found in Appendix (Tables A.14 and A.15).

Table 5
Strategies of the representative participant in the Delay treatment.

Dependent variable:	(1)	(2)	(3)	(4)	(5)
Individual emission in current round					
Trigger: average emission of others > 12					
Lag 1	1.082* (0.496)		0.972* (0.498)		0.919* (0.496)
Lag 2	0.556* (0.324)		0.149 (0.281)		0.142 (0.284)
Any previous round	1.218* (0.630)		1.051* (0.576)		1.053* (0.587)
Trigger: damage > 12 × 39.432					
Lag 1		1.483*** (0.414)	1.118*** (0.337)		1.420*** (0.366)
Lag 2		0.732* (0.375)	0.530 (0.346)		0.548 (0.337)
Any previous round		0.272 (0.573)	−0.315 (0.421)		−0.517 (0.416)
Trigger: personal loss in round payoff					
Lag 1				−0.178 (0.370)	−0.980** (0.442)
Lag 2				0.373 (0.348)	0.015 (0.343)
Any previous round				1.430*** (0.542)	0.758 (0.572)
Constant	6.856*** (0.933)	6.930*** (0.982)	6.811*** (0.925)	6.984*** (0.975)	6.762*** (0.915)
Round dummies	Yes	Yes	Yes	Yes	Yes
Sequence dummies	Yes	Yes	Yes	Yes	Yes
Participant dummies	Yes	Yes	Yes	Yes	Yes
Length past sequence	Yes	Yes	Yes	Yes	Yes
Limited liability	Yes	Yes	Yes	Yes	Yes
Observations	2840	2840	2840	2840	2840
R2	0.5594	0.5552	0.5629	0.5481	0.5643
Adjusted R2	0.5420	0.5376	0.5452	0.5302	0.5461

Notes: Results from OLS regressions are reported. The unit of observation is a participant emission decision in a round. Decisions in round 1 are dropped. Standard errors are clustered at group level.

* $p < 0.1$.
** $p < 0.05$.
*** $p < 0.01$.

A similar pattern emerges from the Delay treatment. Disentangling experiential and observational strategies is statistically more difficult in this design. We limit our focus to just the two rounds following a defection, plus a grim trigger regressor, to minimize the chances of confounding the reaction to actions or to damages (Table 5). In an experiential strategy, a defection occurs if the experienced damage in a round was the outcome of a (previous) average emission in the group above 12, but a robustness check has been performed for other threshold levels (available upon request). Notice that both specifications of experiential strategy in Delay and Stochastic treatments measure an impact on payoffs that is the consequence of both own and others emission choices.

The data suggests that the representative participant increased emissions after experiencing high damages (Table 5, col. 2). The coefficient of the lag 1 regressor for the damage is statistically significant and suggests an immediate reaction, with a smaller coefficient for the lag 2 regressor and an insignificant coefficient for the grim trigger regressor. As already mentioned, the stand-alone estimate of observational strategies using a threshold of 12 yielded statistically significant and positive coefficients for all regressors (Table 5, col. 1). Also in the Delay treatment the reaction to an observed defection seems more permanent than in the Control treatment. When both experiential and observational strategies are jointly estimated, the patterns do not substantially change (Table 5, col. 3). The two classes of strategies are illustrated in Fig. 3 (solid and dashed lines) panel (b). Hence, in the Delay treatment the representative participant is responding with higher emissions both to others’ actions and also to damage higher than a threshold.

Table 6
Strategies of the representative participant in the Stochastic treatment.

Dependent variable:	(1)	(2)	(3)
Individual emission in current round			
Trigger: average emission of others > 12			
Lag 1	1.017*** (0.286)		0.861*** (0.298)
Lag 2	0.351 (0.334)		0.274 (0.296)
Lag 3	-0.211 (0.399)		-0.162 (0.424)
Lag 4	0.136 (0.490)		0.139 (0.515)
Any previous round	0.515** (0.246)		0.443* (0.240)
Trigger: personal accident occurs			
Lag 1		1.447*** (0.296)	1.369*** (0.290)
Lag 2		-0.061 (0.216)	-0.097 (0.221)
Lag 3		-0.058 (0.237)	-0.090 (0.246)
Lag 4		0.156 (0.212)	0.125 (0.206)
Any previous round		-0.088 (0.306)	-0.086 (0.288)
Constant	13.314*** (0.457)	13.283*** (0.516)	12.847*** (0.499)
Round dummies	Yes	Yes	Yes
Sequence dummies	Yes	Yes	Yes
Participant dummies	Yes	Yes	Yes
Length past sequence	Yes	Yes	Yes
Limited liability	Yes	Yes	Yes
Observations	2300	2300	2300
R2	0.3801	0.3918	0.3984
Adjusted R2	0.3525	0.3647	0.3702

Notes: Results from OLS regressions are reported. The unit of observation is a participant emission decision in a round. Decisions in round 1 are dropped. Standard errors are clustered at group level.

* $p < 0.1$.

** $p < 0.05$.

*** $p < 0.01$.

From a behavioral stand point, participants may exhibit discontinuities in strategies when payoffs move into the loss domain. The experimental design allows to study also this possibility, given that in all treatments negative round payoffs were possible. Using the same econometric technique outlined above for strategy estimation, we tracked the response of a representative participant to the experience of a negative round payoffs in the Control and Delay treatments. In the Stochastic treatment the event of an accident coincide with the experience of a negative round payoff, and hence it is impossible to disentangle the effects of each component on strategies. [Result 5](#) below summarizes the findings.

Result 5. (Reaction to losses). *In the Control and Delay treatments, the experience of negative round payoffs modifies the strategy of the representative participant: it reduces the magnitude of the immediate response to a perceived defection.*

Support for [Result 5](#) comes from [Tables 4 and 5](#). The response to losses could share an element of punishment for the high levels of others' emissions causing the loss itself, as well as behavioral elements related to loss aversion and other factors. For this reason, it is important to disentangle the two in the empirical analysis. In the Control treatment, an estimate that tracks the reaction to the experience of negative round payoffs shows a statistically significant increase in emissions in the round following the loss (a positive lag 1 coefficient in [Table 4](#), col. 3). However, when jointly estimating an observational strategy of a trigger type together with the reaction to losses, the net effects are drastically different ([Table 4](#), col. 4): the sign of the

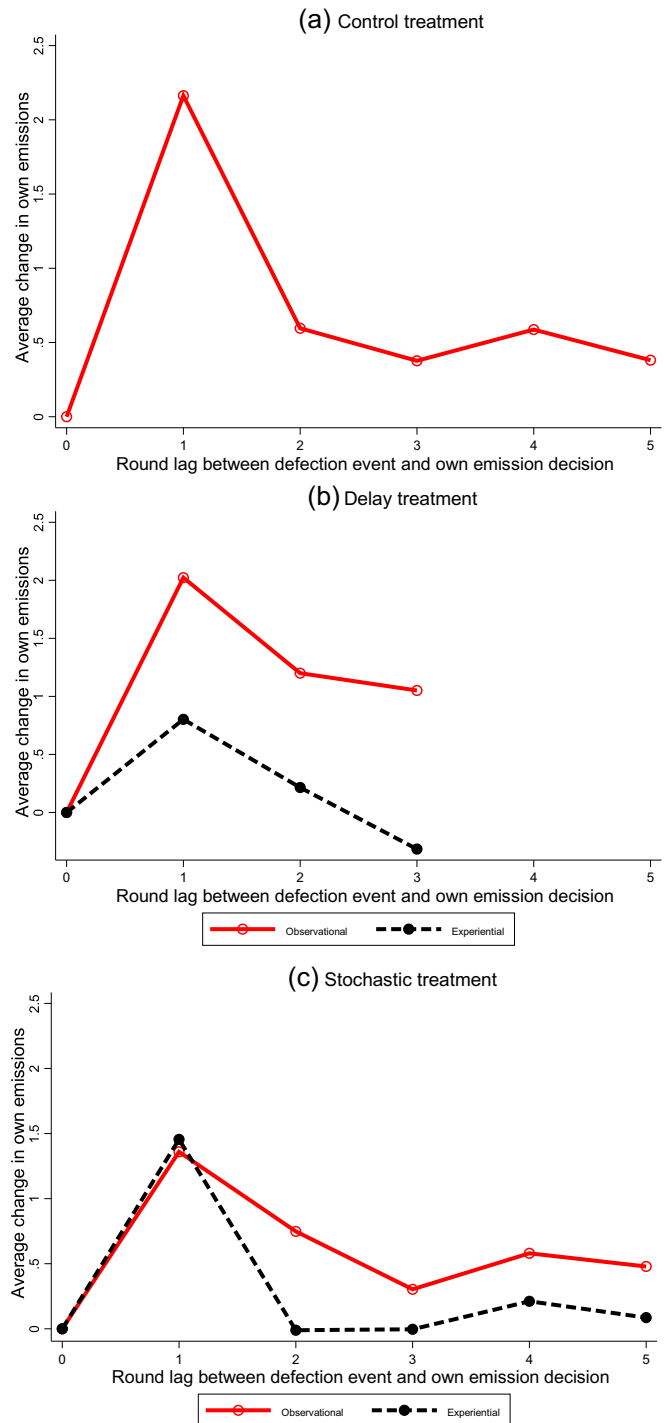


Fig. 3. Strategies of the representative participant. Notes: Own emission change comes from regressions coefficients summing up the estimated coefficient for “Any previous round” and “Lag X”. For Control, see [Table 4](#) col. 2; for Delay, see [Table 5](#) col. 3; for Stochastic, see [Table 6](#) col. 3.

statistically significant coefficient becomes close to zero (Lag 1), and the net effect is negative when summing up the coefficient of “Any previous round” with those of the various lags. When taken together, these two regressions support [Result 5](#) and show that without controlling for the use of a trigger strategy we would have drawn the wrong conclusions about the behavioral effects of a loss. The reason is that the more canonical response due to a punishment for high emissions of others quantitatively dominates the behavioral response

to losses, at least if we focus on the round immediately following the event.

In the Delay treatment the findings are analogous. An estimate that tracks the reaction to the experience of negative round-payoffs shows a permanent increase in emissions (positive coefficient for Any previous round in Table 5, col. 4). However, when jointly estimating an observational strategy of a trigger type together with the reaction to losses, the net effects are drastically different (Table 5, col. 5): we observe a statistically significant negative coefficient for the lag 1 regressor, which remains negative also when summed up with the coefficient of “Any previous round”. Also in this treatment, these two regressions support Result 5.

6.3. Strategies at the individual level

The empirical evidence on strategies in Section 6.2 is compatible with everyone responding to both emissions and damages, and to the presence of two separate types of decision-makers, those who respond exclusively to emissions, and those that respond exclusively to damages. The theoretical and empirical implications of these two scenarios are rather different though, which is why we also carried out a classification of individuals.

In Section 4 we considered a scenario with homogeneous decision-makers in terms of strategy adopted. However, the presence of heterogeneity in behavior may significantly affect emissions. For example, it may require a significant period of learning to envisage other decision-makers' strategies and to build cooperation. One could expect that during this learning process in the Delay and the Stochastic treatments, where experience and observation may be decoupled for some decision-makers, initial emissions are kept cautiously low. At the same time, the learning process may not converge fast enough and the coexistence of experiential and observational decision-makers in the same group may induce spiraling emissions.

We now explain how we classified the participants. The algorithm we used aims at identifying strategies of a “trigger” type where an individual deterministically transitions from a cooperative mode to a punishment mode in the round following an event that is considered a defection. The definition of defection depends on the class of strategy, either experiential or observational, and is associated to a given threshold. The algorithm defines as defection either an observed average action of others above a threshold or the experience of damage, which takes the form of a random accident in the Stochastic treatment, or of a damage level beyond a threshold in the Delay treatment. We check whether each individual's behavior is compatible with an observational trigger strategy and/or an experiential trigger strategy. The unit of observation is a participant in a sequence.

In the Control treatment we cannot identify experiential strategies because there is no decoupling between emission actions and damages. Nonetheless, we can properly identify decision makers following grim trigger or T -round punishment strategies. When basing our counting on these observational trigger strategies, about 34% (75) of the individuals can be classified in this category. In particular, we identify individuals belonging to the observational strategy category by looking at the participant's behavior when making the largest emission increment from one round to the next, $e(t) - e(t - 1)$. To be considered an observational trigger strategy two conditions must be fulfilled. First, this largest emission increment must be in response to an emission increment of the other three group's members (strictly positive on average). Second, when – within a sequence – the participant exhibits multiple instances of largest emission increments, we require that the earliest instance satisfies the first condition above and, in addition, that the same condition applies to the average taken across all instances. This definition applies to all treatments and implicitly captures the idea that a defection occurs if the average emission by the other three members of the group is above a given threshold. This threshold is individual-specific. Some individuals could also follow

strategies other than trigger. The algorithm always places an individual in a sequence of just one or two rounds in the “unclear” category because the data are too sparse. Individuals with constant emissions over time, or emissions that monotonically decline also belong to the “unclear” category.

Similarly, an individual belongs to the category of experiential trigger strategy if her maximal increase in emission between two rounds follows a defection. Here, however, the definition of defection is tied to the personal level of damage. In the Delay treatment, the definition of a defection event follows an analogous rule as in observational strategies but using damages. The threshold that an individual employs to define a defection could be different between experiential and observational strategies.²⁰ Again, we identified the threshold of the experiential strategy by looking at the participant's behavior when making the largest emission increment over two subsequent rounds, $e(t) - e(t - 1)$. To be classified as following an experiential strategy, the individual must have performed this jump in emission in response to a strictly positive damage increment over the previous rounds. In the Stochastic treatment, a defection event occurs every time the individual experiences an accident.²¹

The outcomes of this classification algorithm are illustrated in Fig. 4 and discussed below.

Result 6. (Heterogeneous strategies). *In the Control and Delay treatments, some participants react exclusively to others' actions, some exclusively to changes in payoffs and a third group reacts to both payoffs and actions.*

Among those participants who use trigger strategies, some exclusively respond to actions, others exclusively respond to damages, and another set responds to both actions and damages. These three sets of participants are roughly similar in size. Fig. 4 illustrates that 42%–45% of the classified individuals fall into the observational strategy category in the Delay ($N = 39$) and the Stochastic ($N = 41$) treatments, respectively. 19%–30% into the experiential strategy ($N = 18$ and $N = 27$, respectively). About 39%–25% of classified individuals belong to both categories.

We test whether there are differences in strategy adoption between the Delay and Stochastic treatments using a Probit regression and report no significant effects (Table 7). Although the games and the classification algorithm are in part treatment-specific, we find similar shares of participants who can be classified as observational or as experiential (p -values of Stochastic dummy are $p = 0.476$ and $p = 0.123$, respectively).

No systematic difference between rich and poor emerges in the type of strategy adopted. Instead, the lower is the level of rule understanding about the experiment (variable Mistakes in the quiz), the more likely it is that the participant adopts an experiential strategy. Such regularity does not appear for the adoption of observational strategies, but instead there is a positive and significant effect of the length of the current sequence. To evaluate this evidence one must adjust for the inclusion in Table 7 of all the unclassified individuals. When removing them, we find that the higher is the level of understanding of the rules, the more likely it is that a participant follows an observational strategy (Probit regression, p -value = 0.066, $N = 184$). Moreover, the coefficient of the length of the current sequence loses significance: its effect in Table 7 most likely originates from the fact that participants in longer sequences are easier to classify.

These findings from the classification of individuals reinforce Results 3 and 4 obtained for the representative participants: a large fraction of participants follow a trigger strategy; moreover, some participants

²⁰ With respect to the estimate carried out for the representative participants, in the classification of individuals (i) the strategy thresholds can vary by individual, and (ii) the same individual may adopt a different threshold for observational and experiential strategies.

²¹ Hence, there is no issue of defining a threshold as in the observational strategies. The definition of defection is the same for everyone.

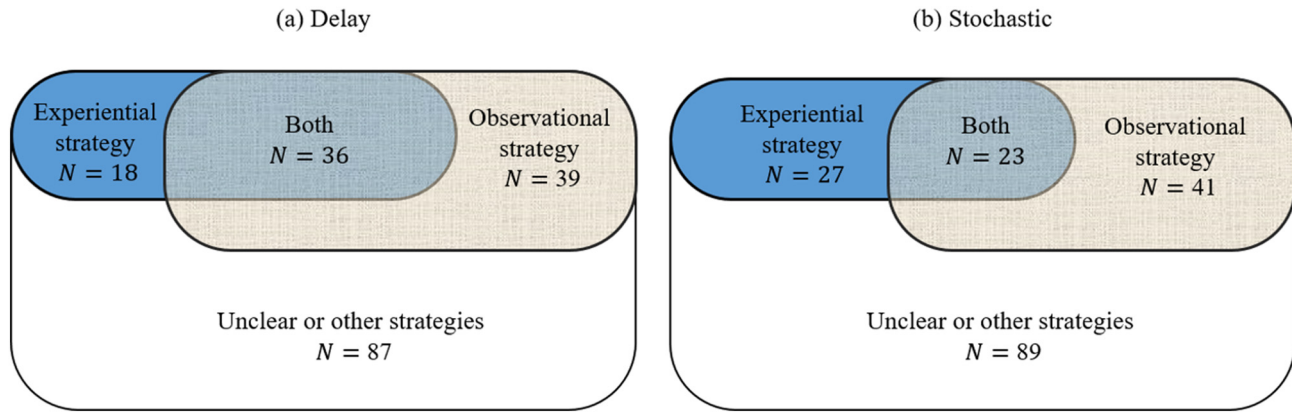


Fig. 4. Classification of participants by strategy employed. Notes: The unit of observation is a participant in a sequence. The number of observations is 180 participants.

respond both to the observed actions as well as to the experienced damage. The corroboration of Result 4 comes from using an empirical technique that is distinct from the regression analysis of Tables 5 and A.14 both in terms of criteria to define strategy types and in terms of unit of analysis. Hence, the classification of individuals strengthens the findings by adding a new point of view on the data.

We conclude our analyses by studying how within-group heterogeneity in strategies affects groups' cooperation. Table 8 reports estimates from an OLS regression where the dependent variable is the average group emission in a sequence and the main covariates are the number of experiential and observational participant in the group. The regression also controls for sequences length and order. We find that a larger number of experiential participants is associated with a higher average group emission. The coefficient is statistically significant in Stochastic. Instead, the number of observational participants does not affect the average group emission in a statically significant way.

7. Discussion and concluding remarks

We experimentally show that two typical features of climate change, namely the delayed and uncertain damages originating from greenhouse gas emissions have relevant and unexpected consequences on mitigation behavior. These features, which are reproduced in our experiment, remove the tight link between the emission action and the consequences in terms of damages. Such link, which characterizes most

Table 7
Determinants of strategy adoption.

Dependent variable:	(1)	(2)	(3)
Individual strategy adopted (0, 1)	Observational	Experiential	Both
Stochastic treatment dummy	0.035 (0.049)	0.054 (0.035)	−0.065 (0.040)
Rich	−0.058 (0.048)	0.010 (0.034)	0.039 (0.040)
Mistakes in the quiz	−0.006 (0.026)	0.040*** (0.015)	0.041** (0.018)
Length of the current sequence	0.007*** (0.002)	0.002 (0.001)	0.004*** (0.001)
Sequence	−0.015 (0.023)	0.019 (0.017)	0.003 (0.021)
Observations	360	360	360

Notes: Marginal effects from Probit regressions are reported. Delay and Stochastic treatments only. The unit of observation is a participant in a sequence. Standard errors are clustered at the level of a participant. The variable "Mistakes in the quiz" counts the number of mistakes made by the participant in the quiz on the instructions.

* $p < 0.1$.
** $p < 0.05$.
*** $p < 0.01$.

studies on cooperation, facilitates learning about others' preferences, level of rationality, and strategies, as well as about the rules of interaction (Bereby-Meyer and Roth, 2006).

Sustaining cooperation on climate issues without a binding international treaty requires informal punishments upon deviations from an agreement. With delayed or uncertain damages, coordinating on a mitigation policy may be harder when some decision-makers condition their actions on emissions and others on actual damages. From a policy point of view it is thus interesting to determine if and when decision-makers decide to react to emissions or to damages. In fact, most theories of long-run cooperation assume homogeneous decision-makers who just focus on actions. When instead decision-makers differ in their attitude to reactions, it is important to know if cooperation is at risk.

Here we design and carry out a laboratory experiment to study the ability of participants to cooperate under different damage functions and with a decoupling of emissions and damages. Clean evidence on these issues is hard to gather with observational data. We report two major sets of results. First, a sizable share of participants is of an experiential type, in the sense of conditioning their emission actions on the level of damages that they individually experience, despite having the possibility to observe also the emission choices of others. Experiential

Table 8
Relation between aggregate emissions and group composition.

Dependent variable:	(1)	(2)
Average emission in a group	Delay	Stochastic
Number of experiential types in the group	0.814 (0.549)	1.155* (0.617)
Number of observational types in the group	−0.747 (0.466)	0.055 (0.377)
Sequence dummies		
Sequence 2	−0.340 (1.180)	0.651 (1.089)
Sequence 3	−1.435 (0.936)	0.312 (0.882)
Sequence 4	0.684 (1.680)	
Length of current sequence	0.046 (0.057)	−0.012 (0.042)
Length of past sequence	−0.146 (0.090)	−0.069* (0.036)
Constant	9.614*** (1.495)	8.560*** (1.104)
R2	0.4251	0.2111
Observations	45	45

Notes: Results from OLS regressions are reported. The unit of observation is a group in a sequence.

* $p < 0.1$.
** $p < 0.05$.
*** $p < 0.01$.

types simply employ personal payoffs as a rule of thumb for decisions. This behavioral finding is novel and contrasts with customary assumptions in theoretical models of cooperation, where everyone does benefit from taking into account the most accurate and timely information available (observational type). A variety of reasons can explain why a decision-maker is experiential or observational. Our evidence suggests that the level of understanding of the situation is systematically worse for experiential than observational types (i.e. types reacting to the observation of emission actions). Experiential decision-makers may find it very costly to keep track and interpret all the information available and thus may fail to appreciate the exact causal connections or the strategies of the others.

Heterogeneity in groups is typical and very likely at the international level when dealing with many countries with widely different political regimes and institutions. Interestingly, in the experiment those groups with more experiential types show significantly higher levels of emissions, a result that may originate from miscoordination. Consider, for example, the interaction between one experiential and one observational decision maker. A stochastic accident would cause the experiential player to increase emissions. The observational player would interpret it as a unilateral deviation, which triggers a punishment. This can ignite a spiral of emissions that unravels cooperation. Similarly, with delayed damages, an observational type may underestimate the reaction of others to a deviation from a cooperation agreement because the experiential type acts later in response to damages. This reasoning may provide an explanation for the increasing time trend of emissions shown in the Delay treatment, the consequence of unfortunate realizations of spiraling reactions. The Stochastic treatment, instead, does not exhibit a time trend. In this treatment, the risk of spiraling reactions is in fact lower because the reaction of experiential types to a deviation by an observational type is smoothed out at the group level due to the asymmetries of the individual accidents. We thus expect that a situation with global common shocks would be more prone to spiraling reactions.

As a second set of results, we find that, overall, the presence of delayed and stochastic damages did not impair cooperation among decision-makers at an aggregate level. Delayed and stochastic damages not only induce some decision makers to act on personally experienced damages. They are also relevant for other dimensions which may reduce the level of emissions. In particular, risk and loss aversion could lead to lower emission with random or delayed damages (the latter because of the uncertain termination date). In light of the higher emissions in groups with several experiential decision-makers, we interpret the aggregate emissions containment under stochastic and delayed damages as the consequence of lower emissions in groups with most observational decision-makers.

With the usual caveats on the limits of external validity of experimental investigations, we believe these findings are relevant for policy design because they call for a careful consideration of how international cooperation emerges and is enforced. In particular, the different attitude towards realized damages of national decision-makers (and their public opinion) may prove a key factor for cooperation. Policy makers may be less concerned with pollutants that display their negative consequences at different times after emission. However, this may be a risk and a missed opportunity to cooperatively and effectively address climate change mitigation.

Appendix A. Supplementary data

Supplementary data to this article can be found online at: <https://doi.org/10.1016/j.eneco.2017.10.029>.

References

- Abreu, D., Pearce, D., Stacchetti, E., 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58 (5), 1041–1063.
- Barrett, S., Dannenberg, A., 2012. Climate negotiations under scientific uncertainty. *Proc. Natl. Acad. Sci. U. S. A.* 109 (43), 17372–17376.
- Barrett, S., Dannenberg, A., 2014. Sensitivity of collective action to uncertainty about climate tipping points. *Nat. Clim. Chang.* 4 (1), 36–39.
- Bereby-Meyer, Y., Roth, A.E., 2006. The speed of learning in noisy games: partial reinforcement and the sustainability of cooperation. *Am. Econ. Rev.* 96 (4), 1029–1042.
- Burke, M., Hsiang, S.M., Miguel, E., 2015. Global non-linear effect of temperature on economic production. *Nature* 527 (7577), 235–239.
- Calzolari, G., Casari, M., Ghidoni, R., 2016. Carbon is forever: a climate change experiment on cooperation. *Quaderni-Working Paper DSE N.* 1065.
- Camera, G., Casari, M., 2009. Cooperation among strangers under the shadow of the future. *Am. Econ. Rev.* 99 (3), 979–1005.
- Dal Bó, P., Fréchette, G.R., 2017. On the determinants of cooperation in infinitely repeated games: a survey. *J. Econ. Lit.* (forthcoming).
- Dannenberg, A., Löschel, A., Paolacci, G., Reif, C., Tavoni, A., 2015. On the provision of public goods with probabilistic and ambiguous thresholds. *Environ. Resour. Econ.* 61 (3), 365–383.
- Dutta, P.K., 1995. A folk theorem for stochastic games. *J. Econ. Theory* 66 (1), 1–32.
- Dutta, P.K., Radner, R., 2004. Self-enforcing climate-change treaties. *Proc. Natl. Acad. Sci. U. S. A.* 101 (14), 5174–5179.
- Falk, A., Heckman, J.J., 2009. Lab experiments are a major source of knowledge in the social sciences. *Science* 326 (5952), 535–538.
- Friedman, J.W., 1971. A non-cooperative equilibrium for supergames. *Rev. Econ. Stud.* 38 (1), 1–12.
- Fudenberg, D., Maskin, E., 1986. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54 (3), 533–554.
- Fudenberg, D., Levine, D., Maskin, E., 1994. The folk theorem with imperfect public information. *Econometrica* 62 (5), 997–1039.
- Fudenberg, D., Rand, D.G., Dreber, A., 2012. Slow to anger and fast to forgive: cooperation in an uncertain world. *Am. Econ. Rev.* 102 (2), 720–749.
- Green, E.J., Porter, R.H., 1984. Noncooperative collusion under imperfect price information. *Econometrica* 52 (1), 87–100.
- IPCC, 2014. Climate change 2014: mitigation of climate change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.
- Karle, H., Kirchsteiger, G., Peitz, M., 2015. Loss aversion and consumption choice: theory and experimental evidence. *Am. Econ. J. Microecon.* 7 (2), 101–120.
- Lauri, S., Holt, C., 2008. Voluntary provision of public goods: experimental results with interior nash equilibria. In: Plott, C., Smith, V. (Eds.), *Handbook of Experimental Economics Results* vol. 1. North-Holland, pp. 792–801.
- Milinski, M., Sommerfeld, R.D., Krambeck, H.-J., Reed, F.A., Marotzke, J., 2008. The collective-risk social dilemma and the prevention of simulated dangerous climate change. *Proc. Natl. Acad. Sci. U. S. A.* 105 (7), 2291–2294.
- Nicklisch, A., Köke, S., Lange, A., 2016. Is adversity a school of wisdom? Experimental Evidence on Cooperative Protection Against Stochastic Losses. Mimeo
- Nordhaus, W.D., 2010. Economic aspects of global warming in a post-Copenhagen environment. *Proc. Natl. Acad. Sci. U. S. A.* 107 (26), 11721–11726.
- Pevnitskaya, S., Ryvkin, D., 2013. Environmental context and termination uncertainty in games with a dynamic public bad. *Environ. Dev. Econ.* 18 (01), 27–49.
- Robbins, D., 1971. Partial reinforcement: a selective review of the alleyway literature since 1960. *Psychol. Bull.* 76 (6), 415.
- Sherstyuk, K., Tarui, N., Ravago, M.-L.V., Saijo, T., 2016. Intergenerational games with dynamic externalities and climate change experiments. *J. Assoc. Environ. Resour. Econ.* 3 (2), 247–281.
- Stern, N., 2015. *Why Are We Waiting? The Logic, Urgency, and Promise of Tackling Climate Change*. MIT Press.
- Tavoni, A., Dannenberg, A., Kallis, G., Löschel, A., 2011. Inequality, communication, and the avoidance of disastrous climate change in a public goods game. *Proc. Natl. Acad. Sci. U. S. A.* 108 (29), 11825–11829.
- Tollefson, J., 2016. Paris climate deal hinges on better carbon accountability. *Nature* 529 (7587), 450.
- Wagner, G., Weitzman, M.L., 2015. *Climate Shock: The Economic Consequences of a Hotter Planet*. Princeton University Press.
- Yamamoto, Y., 2012. Characterizing belief-free review-strategy equilibrium payoffs under conditional independence. *J. Econ. Theory* 147 (5), 1998–2027.