Dipartimento di Economia Internazionale, delle Istituzioni e dello Sviluppo (DISEIS)

WORKING PAPER SERIES

WP 2603

July 2026

Is hostile behavior intuitive or deliberative? A Hawk-Dove experiment with a varying harshness of conflict

Ennio Bilancini, Leonardo Boncinelli, Pablo Marcos-Prieto, Chiara Nardi

Prior to being published, DISEIS working papers are subject to an anonymous refereeing process by members of the editorial board or by members of the department.

DISEIS Working Papers often represent preliminary work which are circulated for discussion and comment purposes. Citation and use of such a paper should account for its provisional character. A revised version may be available directly from the author.

DISEIS WP 2603

DISEIS
Dipartimento di Economia Internazionale, delle Istituzioni e dello Sviluppo
Università Cattolica del Sacro Cuore
Via Necchi 5
20123
Milano

Is hostile behavior intuitive or deliberative? A Hawk-Dove experiment with a varying harshness of conflict

Ennio Bilancini[a], Leonardo Boncinelli[b], Pablo Marcos-Prieto[a], Chiara Nardi[*,c]

[a]IMT School for Advanced Studies Lucca, Italy
[b]University of Florence, Italy
[c]Università Cattolica del Sacro Cuore, Italy

Abstract

Using a one-shot Hawk–Dove game, we experimentally investigate the effect of different cognitive modes—intuitive (induced by Time Pressure), deliberative (by Time Delay), and motivated deliberative (by Time Delay combined with a written motivation)—on the propensity to behave hostilely (i.e., to play Hawk). We also examine whether cognitive modes affect responsiveness to payoff incentives by varying the harshness of conflict. Our results show that intuition significantly increases the likelihood of hostile behavior, while motivated deliberation reduces it. The harshness of conflict does not significantly affect behavior, and we find no evidence that its effect differs across cognitive manipulations. However, when restricting attention to subjects in the pooled delay conditions, the effect of harshness becomes statistically significant, indicating that responsiveness to payoff incentives may require deliberation. Consistently, we find that deliberation increases the likelihood that subjects best respond to their own beliefs.

JEL classification: C72; C90; D91.
Keywords: Cognitive manipulations, Intuition, Deliberation, Motivation, Hawk-Dove game, Harshness of conflict

*Corresponding author. Università Cattolica del Sacro Cuore, Department of International Economics, Institutions and Development, Via Necchi 5, 20123 Milan, Italy. E-mail address: *chiara.nardi@unicatt.it*.

1 Introduction

It is widely recognized that cooperation is a pervasive feature of human societies. Hostile behavior, however, is no less widespread. While a large body of research has investigated the role of cognition in shaping cooperative behavior in social dilemmas (see Capraro, 2024, and references therein), comparatively little attention has been devoted to hostile behavior. This is surprising, as hostile interactions are central to many strategic environments in which individuals compete over scarce resources and face the possibility of costly conflict, often under time constraints. In this study, we aim to fill this gap by assessing whether and how cognition affects hostile behavior in such conflict situations.

Investigating this question in naturally occurring settings is challenging, as both the strategic environment and the underlying cognitive processes are difficult to observe and control. For this reason, we rely on a purposefully designed experiment in which conflict is modeled as a one-shot Hawk-Dove game (Maynard Smith, 1982). The game captures simultaneous interactions in which individuals compete for a resource with positive value. If both players choose Hawk, they engage in an escalated conflict and incur a cost. If both choose Dove, the resource is equally divided. If one player chooses Hawk while the other chooses Dove, the former obtains the entire resource while the latter receives nothing. In this setting, choosing Hawk corresponds to hostile behavior, whereas choosing Dove corresponds to peaceful behavior.

To study the role of cognition within this framework, we manipulate cognitive modes by varying decision times, a standard approach in the literature on dual-process theories of decision-making, which distinguishes between intuitive and deliberative modes of reasoning.[1] Fast, low-cost choices are induced by imposing short response times, whereas slower and more reflective choices are encouraged by requiring longer response times and, in some cases, by further asking participants to provide a written motivation for their decision.

Our design comprises four (between-subject) conditions. In the Control (C) treatment, participants choose between Hawk and Dove without time constraints. In the Time Pressure (TP) treatment, participants are required to decide within 10 seconds, thereby promoting more intuitive responses. In the Time Delay (TD) treatment, participants must wait at least 60 seconds before making a decision, encouraging deliberation. Finally, in the Motivated Delay (MD) treatment, participants must wait at least 60 seconds and use this time to write a brief motivation for their intended choice, thereby inducing a stronger form of deliberation.

In addition to manipulating cognitive modes, we vary the strategic environment by considering different degrees of conflict harshness, defined as the severity of the consequences of an

[1] Alternative methods have been proposed to manipulate cognition, in particular to induce intuitive processing, such as cognitive load (Gilbert et al., 1988), ego depletion (Baumeister et al., 2018), and induction (Liu and Hao, 2011; Shenhav et al., 2012); see Rand et al. (2016) for a meta-analytic comparison. Among these, time-based manipulations are the most commonly employed in the Experimental Economics literature.

escalated fight. Conflict harshness affects the expected payoff of hostile behavior and therefore changes the range of beliefs for which playing Hawk is a best response, with implications for both risk dominance and the mixed-strategy Nash equilibrium of the game.

We implement an additional between-subject condition, orthogonal to the cognitive treatments, by varying the payoff associated with the $(Hawk, Hawk)$ outcome. Specifically, we consider environments in which Hawk is a best response for beliefs that the opponent plays Hawk with probability up to 0.78, and environments in which this threshold is progressively reduced, up to 0.18. This variation allows us to study how cognitive modes interact with the strategic structure of the game, and in particular whether intuition and deliberation differ in their sensitivity to payoff incentives.

Our study relates to two strands of literature. First, we add to the experimental literature on cognition and economic behavior. A growing number of studies have examined how different cognitive modes affect behavior across a variety of strategic settings, including Prisoner's Dilemmas and Public Goods Games (Bilancini et al., 2022a; Rand et al., 2014; Merkel and Lohse, 2019), Dictator Games (Bilancini et al., 2023; Achtziger et al., 2015; Rand et al., 2016), Ultimatum Games (Cappelletti et al., 2011; Grimm and Mengel, 2011; Achtziger et al., 2016), Deception Games (Capraro, 2017), and coordination games such as the Stag Hunt (Belloc et al., 2019; Bilancini et al., 2025). We extend this literature by studying the role of cognitive modes in a canonical conflict game.

Additionally, we contribute to the literature on behavior in the Hawk-Dove game, particularly with regard to the role of the harshness of conflict. While this game has been studied from different theoretical perspectives (Benndorf and Martinez-Martinez, 2017; Bilancini et al., 2022b; Arigapudi and Heller, 2025), experimental evidence remains relatively limited and heterogeneous. To the best of our knowledge, only a few studies have examined behavior in this setting, exploring how it is influenced by population matching protocols (two-population vs. one-population; Oprea et al., 2011), payoff transformations and loss aversion (Feltovich, 2011; Wen et al., 2022), and ex-ante inequality (Benndorf et al., 2016; Bigoni et al., 2026). By combining the Hawk-Dove framework with controlled manipulations of cognitive modes, our study provides new evidence on how individuals behave in conflict situations and how such behavior depends on both cognitive processes and payoff incentives.

Among the existing contributions, the work most closely related to ours is Ordali et al. (2024). In their study, the authors examine the effect of mental fatigue—induced through prolongued exertion of cognitive effort—on behavior in a Hawk-Dove game, and their findings can be interpreted as suggesting that hostile behavior in strategic settings is driven by more intuitive decision-making. However, this interpretation relies on a critical assumption, namely that the induced mental fatigue effectively shifts decision-making toward a more intuitive mode. This step is not guaranteed. While mental fatigue may reduce individuals' ability to engage in deliberation, it may also lead to noisier or more erratic choices rather than to genuinely intuitive responses (Bilancini et al., 2024). Our study improves on this identification by directly manipulating the extent of reflection on decisions, acting on response times and explicitly asking to motivate decisions. By introducing conditions that actively foster deliberation and motivated deliberation, our design allows us to identify how increasing levels of reflection shape hostile behavior.

To anticipate our results, we find that hostile behavior is more likely under Time Pressure and less likely under Motivated Delay relative to the Control condition. The effect of Time Delay is not statistically significant in isolation, but it is directionally consistent with Motivated Delay, and the two delay treatments are not statistically different. When pooled together, the delay conditions provide evidence of a reduction in hostile behavior. We interpret these findings as evidence of intuitive hostility and motivated deliberative restraint, and as suggestive evidence that deliberation more generally reduces hostility. We also find that best-response play is less frequent under Time Pressure and more frequent under both Time Delay and Motivated Delay, indicating that deliberation improves the alignment between beliefs and actions. By contrast, behavior does not respond to payoff incentives: the effect of harshness is negative but not statistically significant, although it becomes more pronounced when the two delay conditions are pooled together. Overall, our results indicate that reduced deliberation leads to more suboptimal (in terms of best response to beliefs) hostile choices, whereas increased deliberation promotes a greater integration of beliefs and payoff-relevant information into decision-making.

These results may appear to be at odds with the literature on the intuitive roots of cooperation, and in particular with the social heuristics hypothesis (Rand et al., 2012).[2] According to this hypothesis, individuals develop heuristics in environments with repeated interactions, where cooperation is often payoff-maximizing due to concerns such as retaliation, reputation, and social exclusion. These heuristics are then carried over to one-shot interactions, leading individuals to cooperate intuitively, whereas deliberation enables them to override these heuristics and choose payoff-maximizing actions (Bear and Rand, 2016). However, it must be noted that the Hawk–Dove game has a markedly different strategic structure from the environments typically considered in the intuitive cooperation literature, such as Public Goods Games and Prisoners' Dilemmas, which may help explain the different patterns of behavior observed in our setting. In this sense, our results suggest that intuitive responses in conflict settings may differ systematically from those observed in cooperative environments.

Our findings are nonetheless consistent with the results in Ordali et al. (2024), as well as with evidence from other coordination games such as the Stag Hunt, where time pressure (Belloc et al., 2019) and delay manipulations (Bilancini et al., 2025) are associated with greater

[2]The social heuristics hypothesis has been the subject of substantial debate, with both supportive (Rand et al., 2016) and critical evidence (Tinghög et al., 2013; Kvarven et al., 2020).

risk-taking, reduced payoff sensitivity, and more suboptimal choices.

The remainder of the paper is organized as follows. Section 2 describes the experimental design, including the parametrization of the Hawk–Dove game, the treatments, the procedures, and the research hypotheses. Section 3 presents our experimental results. Section 4 discusses the results and offers concluding remarks.

2 The experiment

2.1 The Hawk-Dove game and treatments

Our workhorse to study hostile behavior is a one-shot Hawk-Dove game. This is a two-person anti-coordination game in which the players receive an endowment $E > 0$ and compete against each other in order to obtain a resource that has a value $v > 0$ for the winner. Each player simultaneously decides whether to compete hostilely, strategy Hawk (H), or peacefully, strategy Dove (D).[3] If both players behave hostilely (i.e., choose H), they engage in an escalated fight and incur a cost $c > v$.[4] If both players behave peacefully (i.e., choose D), the resource is equally divided between them. If one player behaves hostilely and the other behaves peacefully, the former wins the entire resource, while the latter gets nothing beyond the initial endowment. The players' payoffs, which are common knowledge, are summarized in Figure 1.

Assuming that the players are rational and monetary payoff maximizers, the Hawk-Dove game has two Nash equilibria in pure strategies, (D, H) and (H, D). If we denote the probability that the opponent plays D by α, the expected payoff of playing H is $(1-\alpha)(E+\frac{v-c}{2})+\alpha(E+v)$ while the expected payoff of playing D is $(1-\alpha)E+\alpha(E+\frac{v}{2})$. It follows that an individual is indifferent between the two strategies if and only if α is equal to

$$\alpha^* = 1 - \frac{v}{c}$$

where α^*, the threshold probability for the D strategy in the mixed strategy Nash equilibrium, can be interpreted as a measure of the harshness of conflict (Bilancini et al., 2022b). Notably, when the conflict is harsh (i.e., $\alpha^* > 0.5$), the Dove strategy is risk dominant; conversely, when the conflict is mild ($\alpha^* < 0.5$), the Hawk strategy is risk dominant.

To investigate whether the mode of reasoning affects hostile behavior, decisions in the Hawk-Dove game are collected under four treatments which manipulate the response times.

(i) In the *Control (C) treatment*, participants decide between Hawk and Dove at their own pace (i.e., there are no time constraints).

[3] In the experiment, the Hawk and Dove strategies were neutrally labeled as "Option A" and "Option B", respectively.

[4] Assuming $c < v$ would transform the Hawk-Dove game into a Prisoner's Dilemma.

	H	D
H	$E+\frac{v-c}{2}$, $E+\frac{v-c}{2}$	$E+v$, E
D	E, $E+v$	$E+\frac{v}{2}$, $E+\frac{v}{2}$

Figure 1: The Hawk-Dove game

(ii) In the *Time Pressure (TP) treatment*, which is supposed to trigger intuitive reasoning, participants are asked to make a decision within 10 seconds.

(iii) In the *Time Delay (TD) treatment*, which is supposed to trigger deliberative reasoning, participants are asked to wait at least 60 seconds before making a decision.

(iv) In the *Motivated Delay (MD) treatment*, which is supposed to trigger an even stronger deliberative reasoning compared to TD, participants are asked to wait at least 60 seconds before making a decision and to use this time to provide a short written motivation for the decision they intend to make.

In all treatments, the endowment E and the resource value v amount to 160 and 120 Points, respectively (with 10 Points = £0.05). To assess whether the effect of the mode of cognition (if any) is affected by the harshness of conflict, in each treatment we consider seven values of the cost, $c \in \{136, 160, 200, 240, 280, 360, 440\}$ Points. Table 1 summarizes the payoffs for each cost value and the corresponding level of harshness.

Table 1: Parameters and payoffs

E	v	c	$\pi_i(H,H)$	$\pi_i(H,D)$	$\pi_i(D,H)$	$\pi_i(D,D)$	α^*
160	120	136	152	280	160	220	0.12
160	120	160	140	280	160	220	0.25
160	120	200	120	280	160	220	0.40
160	120	240	100	280	160	220	0.50
160	120	280	80	280	160	220	0.57
160	120	360	40	280	160	220	0.67
160	120	440	0	280	160	220	0.73

2.2 Procedures

The experiment was programmed in oTree (Chen et al., 2016) and conducted online in the second half of July 2024. It was approved by the Joint Ethical Committee of Scuola Normale Superiore and Scuola Superiore Sant'Anna (Italy). The participants were recruited through Prolific (Palan and Schitter, 2018) among the UK adult population. Upon entering the study, they were asked to provide informed consent and to read the instructions (reproduced in Appendix A). Before starting the experiment, subjects had to answer some control questions testing their comprehension of the decision task. The experiment did not start until the participants had answered all the questions correctly.[5] We can, therefore, safely assume that they understood the game.

The implementation of time constraints in the Hawk-Dove game depended on the treatment. In the *TP* treatment, participants had a maximum of 10 seconds to make the choice and a countdown timer was displayed at the top of the decision screen. Participants who exceeded the time limit were redirected to a separate screen that urged them to hurry up and highlighted, in a red box, the extra time taken to reach a decision. Remarkably, the participants were unaware that there was an additional decision screen.[6] In the *TD* treatment, participants had to wait at least 60 seconds before making a decision. During the waiting time, the choice buttons remained hidden to ensure that no response could be submitted before the delay elapsed. The same procedure was applied in the *MD* treatment, with the only difference being that participants were instructed to use the waiting time to write a free-form text motivation of at least 60 characters (a live character count was provided on screen). The choice buttons became visible after 60 seconds, but only if the minimum character requirement had been met. Participants were asked to motivate their choice also in the Control, Time Pressure, and Time Delay treatments. However, in these treatments, the motivation – also subject to the 60-character minimum – was collected on a separate screen after the decision had been submitted.[7]

After making their game choices and providing a written motivation, and without having received any feedback, participants had to report their (first-order) beliefs about the frequency of hostile choices. Beliefs were elicited by asking each participant to guess the percentage of other people who chose the Hawk strategy. Each participant had to indicate an integer number between 0 and 100 by using a slider.[8] We gave participants a financial incentive to report their

[5] In compliance with the Prolific guidelines, participants were given three attempts to answer all the questions correctly. Yet, the vast majority of them (around 80%) successfully completed the comprehension test at the first attempt.

[6] The unanticipated additional decision screen allowed us to retain slow-responding participants in the Time Pressure treatment, thereby avoiding attrition that could have compromised the integrity of random assignment across conditions (Bouwmeester et al., 2017).

[7] We asked a motivation also in the Control, Time Pressure, and Time Delay treatments to keep the time taken to complete the experiment – and thus the hourly compensation – as similar as possible across treatments.

[8] To prevent an anchoring bias, participants had to first click on a bar to reveal the slider and then move the

beliefs accurately. They received 40 Points for a correct prediction, 30 Points if their prediction was incorrect by ± 1 percentage point, 20 Points if it was incorrect by ± 2 or 3 percentage points, 10 Points if it was incorrect by ± 4 or 5 percentage points, and 0 Points otherwise. Incentives in the belief task were kept small relative to incentives in the Hawk-Dove game to avoid hedging (Blanco et al., 2010). When participants made their game decisions, they were unaware of the subsequent belief elicitation task. This avoids any influence of beliefs on game decisions.

Finally, participants filled out a post-experimental questionnaire asking them about their (i) demographic characteristics (age and gender), (ii) risk tolerance and (iii) prosocial traits (positive and negative reciprocity, altruism, and trust).[9] The risk tolerance was measured with a non-incentivized question from the German Socio-Economic Panel asking participants to rate their willingness to take risks in general on an 11-point scale ranging from 0 (not at all willing to take risks) to 10 (very willing to take risks). The behavioral validity of this survey risk measure has been confirmed by Dohmen et al. (2011). Prosocial traits were elicited with questions from the Global Preference Survey (Falk et al., 2018, 2023). More specifically, positive reciprocity was measured using two questions. The first question asked participants how much to spend on a thank-you present (from 0 to 30, in increments of 5) in return for the help received from a stranger in reaching a desired destination in an unfamiliar area, while the second question asked participants to self-assess their willingness to return a favor on an 11-point scale from 0 to 10. Negative reciprocity was measured by asking participants to self-assess their willingness (a) to take revenge if they are treated very unjustly even if doing so is costly and (b) to punish someone for unfair behavior, either toward themselves or toward a third person. Altruism was elicited with two questions related to donations. The first question asked participants to rate, again on a 0-10 scale, their willingness to give to good causes without expecting anything in return. The second question was a hypothetical Dictator game in which participants had to indicate how much of £1000 to donate. Finally, trust was measured by asking participants whether they assume that other people only have the best intentions, on a 0-10 scale.

We used a between-subjects design, i.e., each subject was exposed to only one of the four treatments and experienced only one harshness level. Averaging over all treatments, mean earnings amounted to £1.57 (inclusive of a £0.50 fixed participation fee) and participants took less than 8 minutes to complete the experiment. The incentives in the experiment were thus substantial and comparable to the typical hourly compensation offered in laboratory experiments (namely, £12).

pointer to the desired number.

[9] In addition to the demographic characteristics collected in the post-experimental questionnaire (reproduced in Appendix A), we retrieved information on the participants' level of education from Prolific. Prior to the administration of the post-experimental questionnaire, participants provided their personal and social norms, which are analyzed in a separate paper.

Overall, 1120 subjects (62% female, mean age = 37.1 years) participated in the experiment, i.e., 280 participants per treatment. In each treatment, participants were equally divided across the seven harshness levels (i.e., 40 participants per level).

2.3 Hypotheses

The research hypotheses, which we pre-registered in the OSF Registry together with the sample size and data analysis plan (`https://osf.io/q8kra`), concern how cognitive manipulations—namely, Time Pressure, Time Delay, and Motivated Delay—and the harshness of conflict independently and through their interaction affect the frequency of Hawk choices.

We begin by considering the effects of cognitive manipulations. Decision-making under Time Pressure limits the opportunity to reflect on the strategic consequences of one's action. In the Hawk–Dove game, choosing Hawk entails entering a conflict that can lead to unfavorable outcomes if matched with another Hawk. One plausible mechanism is that Time Pressure reduces attention to such contingent outcomes and favors more immediate or salient responses. In addition, prior evidence suggests that, in gain domains, intuitive decision-making may be associated with increased risk-taking (Kocher et al., 2013; Young et al., 2012; Saqib and Chan, 2015). In the present setting, choosing Hawk exposes the subject to conflict risk; thus, both reduced reflection and increased risk-taking would predict more frequent Hawk choices under Time Pressure.

By contrast, imposing a waiting time allows subjects to reflect on the potential consequences of conflict and on the strategic environment. Deliberation may increase the salience of mutually harmful outcomes and encourage more cautious behavior, making subjects less inclined to choose the hostile action.

The Motivated Delay treatment further requires subjects to use the waiting time to articulate the reasoning behind their intended choice. This requirement may induce deeper processing by making payoff consequences and strategic considerations more explicit, thereby strengthening deliberation relative to time delay alone.

Taken together, these considerations lead to the following hypotheses:

Hypothesis 1 (Intuitive hostility): *In the Time Pressure treatment, the frequency of Hawk choices is higher than in the Control treatment.*

Hypothesis 2 (Deliberative kindness): *In the Time Delay treatment, the frequency of Hawk choices is lower than in the Control treatment.*

Hypothesis 3 (Motivated deliberative kindness): *In the Motivated Delay treatment, the frequency of Hawk choices is lower than in the Control treatment.*

Hypothesis 4 (Motivated deliberation): *In the Motivated Delay treatment, the frequency of Hawk choices is lower than in the Time Delay treatment.*

We next turn to the effect of harshness. As harshness increases, the cost of entering the conflict becomes larger, reducing the attractiveness of choosing Hawk. Hence, in all treatments, subjects who respond to payoff incentives would be expected to choose Hawk less frequently when harshness is higher. Accordingly, we state the following hypothesis:[10]

Hypothesis 5 (Equilibrium rationality): *The frequency of Hawk choices decreases with the harshness of conflict.*

Finally, we consider the interaction between the mode of cognition and harshness, as the former may affect not only the level of Hawk play but also the sensitivity of choices to payoff differences. Under Time Pressure, subjects may have limited ability to process the magnitude of payoff differences, leading to reduced sensitivity to changes in harshness. By contrast, under Time Delay and Motivated Delay, subjects may be better able to process payoff information and incorporate it into their decisions. Hence:[11]

Hypothesis 6 (Reduced payoff sensitivity): *The negative relationship between harshness and the frequency of Hawk choices is weaker (i.e., less negative) in the Time Pressure treatment than in the Time Delay and Motivated Delay treatments.*

3 Results

In this section, we describe the data and present the results obtained using the full sample. In Appendix B, we report the results obtained when excluding the non-compliant participants, i.e., those who did not choose within 10 seconds in the Time Pressure treatment.

Before presenting the main results of the study, we verify that our treatments successfully manipulated decision times. Table 2 reports descriptive statistics of the amount of time (in seconds) subjects took to make a choice in the Hawk-Dove game. Our cognitive manipulations worked as expected: participants in Time Pressure responded faster than those in the Control (7.11 vs 15.14 seconds), while participants in Time Delay and Motivated Delay took longer (80.81 and 89.55 seconds, respectively). A series of Wilcoxon rank-sum tests confirms that all pairwise differences in response times between treatments are statistically significant (all p-values < 0.001).[12] We can therefore follow up by presenting all results under the assumption that our treatments generated the intended variation in response times.

[10] In the pre-registration, Hypothesis 5 is stated with reference to the Control treatment. We present it here in a more general form for expositional clarity, as standard game-theoretic reasoning implies that the relationship between harshness and behavior applies across treatments.

[11] Hypothesis 6 has been slightly rephrased relative to the pre-registration to better align with its empirical test, without altering its substantive content.

[12] All p-values in the paper are two-tailed.

Table 2: Mean (and standard deviation) of decision times (in seconds)

Control	Time Pressure	Time Delay	Motivated Delay
15.14	7.11	80.81	89.55
(13.89)	(4.22)	(127.24)	(39.81)

Figure 2: Proportion of Hawk choices by treatment

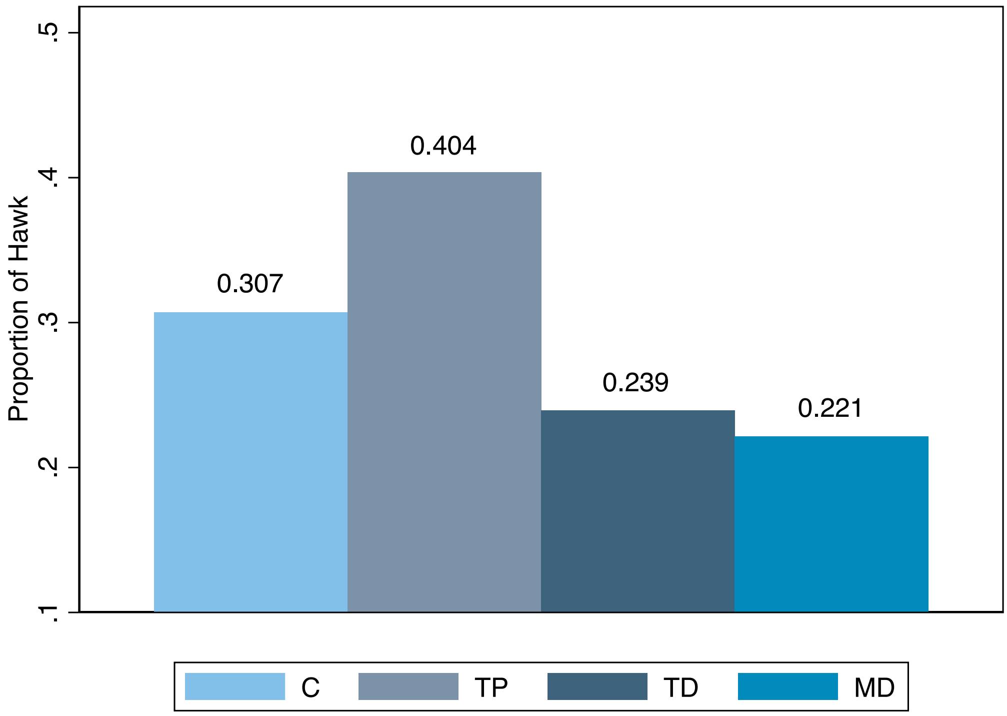

3.1 Descriptive and nonparametric analyses

Figure 2 shows the frequency of Hawk choices across treatments. Its visual inspection reveals three noteworthy features. First, people behave more hostilely under intuition. In the Control treatment, 30.71% of participants choose Hawk, while this percentage increases to 40.36% in the Time Pressure treatment, and the difference is statistically significant (χ^2 test, p-value=0.017). Second, under deliberation, people behave in a less hostile way. Compared to the Control, the proportions of Hawk choices are lower in both the Time Delay treatment (30.71% vs. 23.93%; χ^2 test, p-value=0.072) and the Motivated Delay treatment (30.71% vs. 22.14%; χ^2 test, p-value=0.021). Third, the proportions of Hawk choices do not differ across the Time Delay and Motivated Delay treatments (23.94% vs. 22.14%; χ^2 test, p-value=0.616). In the next subsection, we will employ regression analysis to thoroughly test the predictions presented in Section 2.3.

3.2 Effects of cognitive manipulations and harshness on hostile behavior

Table 3: Probit regressions examining the probability of behaving hostilely

	(1)	(2)
Time Pressure	0.261**	0.289**
	(0.109)	(0.114)
Time Delay	-0.207*	-0.131
	(0.114)	(0.118)
Motivated Delay	-0.266**	-0.247**
	(0.115)	(0.118)
Harshness	-0.377*	-0.371*
	(0.196)	(0.203)
Age		0.003
		(0.003)
Female		-0.015
		(0.088)
Education (in years)		-0.029*
		(0.016)
Risk tolerance		0.303***
		(0.047)
Positive Reciprocity		-0.104*
		(0.057)
Negative Reciprocity		0.050
		(0.048)
Altruism		0.139***
		(0.053)
Trust		0.039
		(0.043)
Constant	-0.331***	-0.094
	(0.118)	(0.307)
Coefficient tests (p-values)		
Time Pressure = Time Delay	< 0.001	<0.001
Time Pressure = Motivated Delay	< 0.001	< 0.001
Time Delay = Motivated Delay	0.615	0.344
Observations	1120	1104
Pseudo R^2	0.023	0.075

Robust standard errors in parentheses.
* p-value < 0.10, ** p-value < 0.05, *** p-value < 0.01.

Table 3 displays the results of probit models aimed at examining the propensity of participants to behave hostilely (i.e., choose Hawk) in the different treatments and the determinants of their behavior. The coefficient of "Time Pressure" in column (1) is positive and significant. Conversely, the coefficient of "Motivated Delay" is negative and significant. We thus confirm that participants are more (less) likely to behave hostilely under intuition (deliberation).[13] Additionally, the t-tests reported at the bottom of the table reveal that the coefficient of "Motivated

[13]Interestingly, we find that the effect of Time Pressure on hostile behavior is mainly driven by males, whereas the effect of Motivated Delay is mainly driven by females. For an experimental study designed to examine whether gender differences exist in how modes of cognition influence hostile behavior, see Bilancini et al. (UnpR).

Delay" is not significantly lower than the coefficient of "Time Delay", providing no evidence that requiring a written motivation produces an additional reduction in hostile behavior beyond that induced by delay alone. These results hold true even if we add controls for participants' demographics and preferences (see column (2)). Among the added control variables, "Risk tolerance" and "Altruism" have an impact on the likelihood of behaving hostilely. In accordance with Hypotheses 1 (intuitive hostility) and 3 (motivated deliberative kindness) and not with Hypotheses 2 (deliberative kindness) and 4 (motivated deliberation), we state:

Result 1. Inducing people to rely on intuition (motivated deliberation) significantly increases (decreases) the likelihood that they behave hostilely.

Result 2. The likelihood of behaving hostilely is not significantly lower under motivated deliberation than under unmotivated deliberation.

We now turn to the analysis of the severity of the conflict. The coefficient of "Harshness" in Table 3 is negative both when excluding and including the control variables. While this is consistent with Hypothesis 5, the effect does not reach statistical significance at conventional levels (5%).

To further investigate whether there are differences in the effect of harshness on hostile behavior by treatment, we estimate probit models in which the interactions between the treatment dummies and the harshness level are added to the set of explanatory variables. The results of this exercise are reported in Table 4 (the reference group is the group of participants assigned to the Control treatment with the lowest harshness level). The coefficient of "Harshness" is negative and insignificant. Moreover, none of the added interaction terms are significant, and the coefficient tests reported at the bottom of Table 4 fail to reject their equality across treatments. Hence, we can conclude that the harshness level does not affect the probability of behaving hostilely and this does not vary across treatments.[14] In line with these findings, and contrary to Hypotheses 5 (equilibrium rationality) and 6 (reduced payoff sensitivity), we state:

Result 3. The harshness of conflict does not significantly influence the likelihood of behaving hostilely, nor does its effect differ across cognitive manipulations.

This null result on harshness might be explained from the fact that, when pooling subjects in the Control and Time Pressure treatments (hereafter, non-delayed conditions), behavior in the game does not depend on harshness (χ^2 test, p-value=0.8567). In contrast, when pooling subjects in the Time Delay and Motivated Delay treatments (delayed conditions), harshness

[14]In Appendix B, we examine the effect of harshness on hostile behavior by estimating the relationship separately within each treatment. The results in the first two columns of Table B3 show that the harshness level does not significantly affect behavior in the Control treatment, both when excluding and when including control variables. These specifications provide a formal test of Hypothesis 5, as originally formulated in the pre-registration.

Table 4: Probit regressions examining the role of harshness in behaving hostilely

	(1)	(2)
Time Pressure	0.147	0.188
	(0.269)	(0.283)
Time Delay	0.042	0.086
	(0.274)	(0.284)
Motivated Delay	-0.007	0.017
	(0.280)	(0.294)
Harshness	-0.183	-0.188
	(0.384)	(0.400)
Time Pressure × Harshness	0.245	0.221
	(0.533)	(0.559)
Time Delay × Harshness	-0.549	-0.482
	(0.552)	(0.568)
Motivated Delay × Harshness	-0.573	-0.589
	(0.566)	(0.593)
Age		0.003
		(0.003)
Female		-0.013
		(0.088)
Education (in years)		-0.029*
		(0.016)
Risk tolerance		0.300***
		(0.047)
Positive Reciprocity		-0.103*
		(0.057)
Negative Reciprocity		0.047
		(0.048)
Altruism		0.143***
		(0.053)
Trust		0.039
		(0.043)
Constant	-0.420**	-0.171
	(0.193)	(0.352)
Coefficient tests (p-values)		
Time Pressure × Harshness = Time Delay × Harshness	0.145	0.210
Time Pressure × Harshness = Motivated Delay × Harshness	0.135	0.152
Time Delay × Harshness = Motivated Delay × Harshness	0.967	0.856
Observations	1120	1104
Pseudo R^2	0.025	0.077

Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

significanlty influences behavior (χ^2 test, p-value=0.0337). This suggests that equilibrium rationality may require deliberation. Figure B2 in Appendix B shows the harshness levels for non-delayed and delayed conditions, while Table B4 reports the corresponding probit regressions.

3.3 Beliefs and best-reply analysis

In this subsection, we present an additional analysis of participants' first-order beliefs, which was not pre-registered. Table 5 reports OLS regressions on the expected percentage of Hawk choices (ranging from 0 to 100). In columns (1) and (2), we employ the same specifications as in Table 3, while in columns (3) and (4), we employ the same specifications as in Table 4. Both when excluding (column (1)) and including (column (2)) control variables, the coefficient for "Time Pressure" is positive and highly statistically significant, indicating that participants under intuition (correctly) expect a higher percentage of Hawk choices compared to those in the Control treatment. The coefficient of "Harshness" in column (2) is instead negative and statistically significant, implying that participants expect fewer hostile choices as the harshness of conflict increases. These effects disappear in the last two columns of Table 5, where interactions terms between treatment dummies and the harshness level are included among the predictors.

A natural question that arises is whether our cognitive manipulations impact the extent to which subjects are able to best respond to their own beliefs. We denote by $\theta_i(\alpha_i^*)$ the belief of player i, facing harshness level α_i^*, about the proportion of players choosing strategy H. Then, player i's *best response correspondence* conditional on beliefs, which we denote by $b_i(\theta_i) \in \{(H, D)\}$, is as follows:

$$b_i(\theta_i) = \begin{cases} H & \text{if } \theta_i(\alpha_i^*) < (1 - \alpha_i^*) \\ \emptyset & \text{if } \theta_i(\alpha_i^*) = (1 - \alpha_i^*) \\ D & \text{if } \theta_i(\alpha_i^*) > (1 - \alpha_i^*) \end{cases}$$

Panel (a) of Figure 3 presents the proportion of best-replying subjects across the different treatments. While there appears to be a trend toward a higher best-reply rate among subjects under deliberation, no significant difference is detected across treatments (χ^2 test, p-value=0.211). Since the Control and Time Pressure treatments show virtually no difference in best-reply rates (χ^2 test, p-value=0.927), we subsequently pool our treatments into two categories: non-delayed (Control and Time Pressure) and delayed (Time Delay and Motivated Delay). Panel (b) of Figure 3 shows a significant increase in the best-reply rate under delay (χ^2 test, p-value=0.043). However, given that the strategy set and best-reply correspondence are discrete, while subjects' beliefs are (approximately) continuous, this analysis may not fully capture the nuances of the relationship between beliefs, best-reply behavior, and harshness. To shed some light on this

Table 5: OLS regressions examining first-order beliefs

	(1)	(2)	(3)	(4)
Time Pressure	5.596***	6.281***	1.913	2.439
	(2.128)	(2.117)	(5.252)	(5.254)
Time Delay	0.204	1.250	-1.199	-0.831
	(2.006)	(2.009)	(4.961)	(4.980)
Motivated Delay	0.232	0.616	-0.932	-1.622
	(1.977)	(1.974)	(5.032)	(5.052)
Harshness	-7.261	-6.999**	-10.644	-11.447
	(3.465)	(3.458)	(7.527)	(7.558)
Time Pressure × Harshness			7.975	8.334
			(10.347)	(10.381)
Time Delay × Harshness			3.037	4.521
			(9.907)	(9.929)
Motivated Delay × Harshness			2.521	4.842
			(10.039)	(10.011)
Age		-0.083		-0.085
		(0.056)		(0.056)
Female		0.396		0.411
		(1.506)		(1.507)
Education (in years)		-0.176		-0.186
		(0.273)		(0.273)
Risk tolerance		3.303***		3.311***
		(0.751)		(0.751)
Positive Reciprocity		-0.673		-0.694
		(0.984)		(0.985)
Negative Reciprocity		1.016		1.006
		(0.847)		(0.849)
Altruism		1.455		1.419
		(0.894)		(0.902)
Trust		-0.051		-0.010
		(0.760)		(0.765)
Constant	48.521***	53.187***	50.084***	55.426***
	(2.182)	(5.353)	(3.769)	(6.153)
Observations	1120	1104	1120	1104
R^2	0.014	0.044	0.014	0.045

Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3: Best reply rate of subjects, conditional on their beliefs, and expected losses and gains

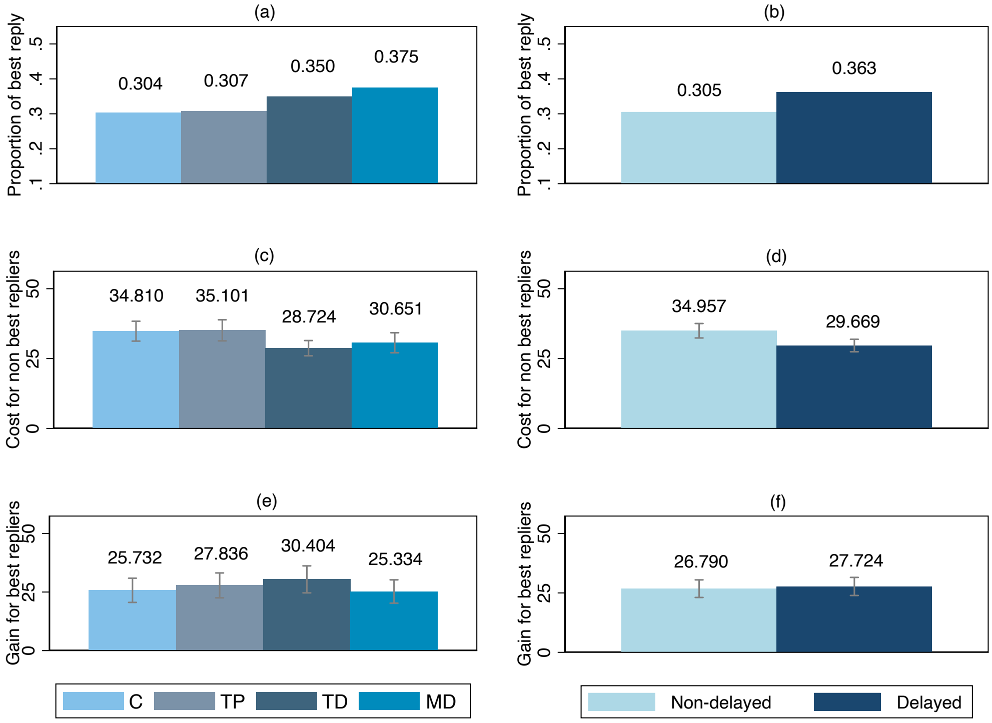

issue, we define the expected cost of not best-replying for player i that does not best reply as:

$$l_i(\theta_i) = \begin{cases} \mathbb{E}[\pi(H)] - \mathbb{E}[\pi(D)] & \text{if } b_i(\theta_i) = H \\ \mathbb{E}[\pi(D)] - \mathbb{E}[\pi(H)] & \text{if } b_i(\theta_i) = D \end{cases}$$

where $\mathbb{E}[\pi(H)]$ is the expected payoff of playing H, while $\mathbb{E}[\pi(D)]$ is the expected payoff of playing D. Equivalently, the expected profit of best-replying for player i that does best reply is simply $g_i(\theta_i) = l_i(\theta_i)$, since the expression is equally constructed but only for those subjects that best-reply. Panel (d) of Figure 3 shows that subjects under delay have a significantly lower expected cost whenever they do not manage to best reply (χ^2 test, p-value=0.008). Overall, subjects under delay are able to best-reply more often, and they fail do so only when the payoff structure is such that they are closer to the point of indifference between both strategies, since their expected cost is lower. Conversely, as presented in Panel (f) of Figure 3, they do not have further gains over subjects not under delay whenever they best-reply (χ^2 test, p-value=0.871), which is compensated for the fact that they do so more frequently. We maintain the disaggregation by treatments in Panels (c) and (e) of Figure 3 for descriptive purposes.

If we plot the proportion of best replies for different harshness levels, it seems that subjects are better able to best reply to their own beliefs for more extreme harshness levels, as seen in Figure B3 in Appendix B. This intuition is confirmed statistically if we test for the proportion of best replying participants for the two maximum and minimum harshness levels against the five central ones (χ^2 test, p-value<0.001), a result that is robust to considering instead the four maximum and minimum harshness levels against the three central ones (χ^2 test, p-value=0.0019). This effect is weakened when we consider only non-delayed subjects (χ^2 test, p-value=0.0137 for five central levels and p-value=0.0565 for three central levels), being driven mainly from subjects in delayed treatment (χ^2 test, p-value=0.0018 for five central levels and p-value=0.0129 for three central levels), as seen in Figure B4 in Appendix B.

4 Conclusions

In this paper, we investigate how cognitive manipulations and the harshness of conflict independently and jointly affect hostile behavior in a one-shot Hawk–Dove game. We manipulate cognitive modes through Time Pressure, Time Delay, and Motivated Delay, and test whether these conditions affect both the level of hostile behavior and its sensitivity to payoff incentives. In particular, we examine whether intuition increases the propensity to choose Hawk, whether deliberation reduces it, whether higher harshness discourages hostile behavior, and whether cognitive modes affect responsiveness to harshness.

Our results provide clear evidence that cognitive modes play a central role in shaping behavior. In line with our hypotheses, we find that Time Pressure significantly increases the likelihood of choosing Hawk relative to the Control treatment, while Motivated Delay reduces it. The effect of Time Delay is not statistically significant in isolation, but it is consistent with that of Motivated Delay, and the two are not statistically different. This absence of a statistically significant difference means that the reduction in hostility is not specific to the motivated condition, but reflects a more general effect of deliberation. Remarkably, when pooled together, the delay conditions provide evidence of a reduction in hostile behavior relative to the Control treatment (χ^2 test, p-value=0.016), suggesting that the difference is quantitative rather than qualitative and may reflect limited statistical power rather than the absence of an underlying effect. Overall, the delay treatments consistently point to a reduction in hostility under deliberation.

While standard game-theoretic reasoning predicts that higher harshness should reduce the attractiveness of the Hawk strategy, we find only limited support for the role of payoff incentives. Although the estimated effect of harshness is negative, it does not reach statistical significance at conventional levels, and we find no evidence that sensitivity to harshness differs across cognitive conditions. Yet, when restricting attention to subjects in the pooled delay conditions, the effect of harshness becomes statistically significant, indicating that responsiveness to payoff incentives may require deliberation.

Examining first-order beliefs and best-reply behavior provides further insight into the underlying mechanisms. On the one hand, subjects under Time Pressure expect a higher frequency of hostile behavior. However, this is not accompanied by a corresponding adjustment in their actions, indicating that beliefs and behavior are not strategically aligned under intuition. On the other hand, deliberation does not substantially alter beliefs, but improves the extent to which actions are consistent with best responses to those beliefs. In particular, subjects under delay conditions best respond more frequently to their beliefs and incur lower expected costs when they fail to do so. This suggests that deliberation enhances the integration of beliefs and payoff-relevant information into decision-making.

Two aspects of the empirical analysis warrant further discussion. First, in line with our pre-registration, we adopt an intention-to-treat approach and include the full sample in the analysis. A closer inspection of the data reveals that the effect of Time Pressure is driven by non-compliant subjects. While this could raise concerns about the robustness of the result, we do not consider it a major issue. Non-compliant subjects remain significantly faster than those in the Control treatment (Wilcoxon rank-sum test, p-value = 0.022), indicating that they are still exposed to the time constraint. Moreover, non-compliance is likely to reflect difficulty in coping with the time limitation rather than inattention or misunderstanding. In this sense, these subjects may be precisely those for whom the manipulation is most binding. If the effect of Time Pressure operates through perceived cognitive constraint rather than strict compliance with the time limit, stronger effects are expected among individuals who struggle to respond within the allotted time.

Second, the weak effect of harshness may reflect limited salience of payoff differences in a one-shot environment or potential non-linearities in behavioral responses. These factors require cognitive effort to be processed, which may explain why sensitivity to harshness emerges only under deliberation. Future research could further investigate this interpretation by experimentally varying both cognitive mode and the salience and range of harshness levels. In particular, varying harshness within subjects—while maintaining one-shot incentives in each decision—would allow participants to face multiple levels of conflict severity and test more directly whether they adjust their behavior to changes in payoff incentives, see Bilancini et al. (UnpR). Moreover, comparing such designs with repeated Hawk–Dove games that provide feedback would help disentangle the roles of immediate payoff processing and learning in shaping sensitivity to harshness. Finally, using a finer grid of harshness levels would allow researchers to detect potential non-linear responses and identify the regions in which payoff differences have the greatest behavioral impact.

Our findings have important implications for real-world conflict situations. In environments such as negotiations, organizational disputes, or even large-scale conflicts, decisions are often made under time pressure. Our results suggest that such conditions may foster more hostile behavior, while deliberation and reflection may help mitigate conflict. For instance, in high-stakes settings such as diplomatic or military decision-making, introducing mechanisms that slow down decision processes or encourage reflection may reduce the likelihood of escalation. More generally, our findings highlight the importance of considering not only incentives, but also the cognitive conditions under which decisions are made.

Appendix A: Experimental instructions, belief elicitation and questionnaire

Instructions

General

In this study, you will be paired with another person. In the following we will refer to this person as the Other. You will never learn the identity of the Other and the Other will never learn yours. All amounts will be expressed in Points rather than pounds. The exchange rate is 10 Points = £0.05.

Your task

You, as well as the Other, have to choose one of two options, which will be labeled A and B, by clicking on the corresponding button. [*Participants in the Time Pressure treatment read*: You will have 10 seconds to make your choice.] [*Participants in the Time Delay treatment read*: You will have to wait at least 60 seconds before making your choice, i.e., the choice buttons will appear once the waiting period is over.] [*Participants in the Motivated Delay treatment read*: Before selecting an option, you will have to motivate your choice by writing one or more sentences (60 characters at least) in 60 seconds or more.]

Your earnings

The Points you earn will depend on the choice you make and the choice made by the Other as displayed in the following table:

	Points YOU earn	Points the OTHER earns
You choose **A** and the Other chooses **A**	x	x
You choose **B** and the Other chooses **A**	160	280
You choose **A** and the Other chooses **B**	280	160
You choose **B** and the Other chooses **B**	220	220

The actual value of x will be shown in the decision screen.

Belief elicitation task

(We ask you to guess the percentage of participants (excluding yourself) who choose option A in the decision task. To indicate your guess, you have to choose an integer number between 0

and 100 by clicking on the slider below and moving the pointer to the desired number. For your guess, you can earn extra Points, which will be added to your earnings from the decision task. The Points you earn will depend on how close your guess is to the actual (rounded) percentage of participants who choose option A. If your guess is exactly right, you earn 40 Points. If your guess is incorrect by 1 percentage point, you earn 30 Points. If your guess is incorrect by 2 or 3 percentage points, you earn 20 Points. If your guess is incorrect by 4 or 5 percentage points, you earn 10 Points. If your guess is incorrect by 6 or more percentage points, you earn 0 Points.)

In your best assessment, what is the percentage of other participants who choose A in the decision task?

0 100

Post-experiment questionnaire

(We kindly ask you to answer a short questionnaire. Your responses are completely confidential.)

1. What is your age?
2. With which gender do you identify?
   - Male
   - Female
   - Other
3. Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Please indicate your answer on a scale from 0 to 10, where 0 means “unwilling to take risks” and 10 means “fully prepared to take risks”.
4. Please think about what you would do in the following situations.

   You are in an area you are not familiar with, and you realize you lost your way. You ask a stranger for directions. The stranger offers to take you to your destination. Helping you costs the stranger about £20 in total. However, the stranger says he or she does not want any money from you. You have six presents with you. The cheapest present costs £5, the most expensive one costs £30. Do you give one of the presents to the stranger as a “thank-you”-gift? If so, which present do you give to the stranger?

   - No present
   - The present worth £5
   - The present worth £10
   - The present worth £15
   - The present worth £20
   - The present worth £25
   - The present worth £30

Today you unexpectedly received £1,000. How much of this amount would you donate to a good cause? (Values between 0 and 1000 are allowed.)

5. How well do the following statements describe you as a person? Please indicate your answer on a scale from 0 to 10, where 0 means "does not describe me at all" and a 10 means "describes me perfectly".
   - When someone does me a favor I am willing to return it.
   - I assume that people have only the best intentions.
   - If I am treated very unjustly, I will take revenge at the first occasion, even if there is a cost to do so.
   - I assume that people have only the best intentions.
6. We now ask for your willingness to act in a certain way in three different areas. Please indicate your answer on a scale from 0 to 10, where 0 means "completely unwilling to do so" and 10 means "very willing to do so".
   - How willing are you to punish someone who treats you unfairly, even if there may be costs for you?
   - How willing are you to punish someone who treats others unfairly, even if there may be costs for you?
   - How willing are you to give to good causes without expecting anything in return?

Appendix B: Supplementary data analysis

Results when excluding non-compliants

Table B1: Probit regressions examining the probability of behaving hostilely (compliants only)

	(1)	(2)	(3)	(4)
Time Pressure	0.108	0.150	0.154	0.182
	(0.116)	(0.121)	(0.279)	(0.294)
Time Delay	-0.207*	-0.133	0.042	0.078
	(0.114)	(0.118)	(0.275)	(0.284)
Motivated Delay	-0.267**	-0.252**	-0.007	0.005
	(0.115)	(0.118)	(0.280)	(0.295)
Harshness	-0.481**	-0.474**	-0.183	-0.205
	(0.200)	(0.208)	(0.384)	(0.400)
Time Pressure × Harshness			-0.095	-0.064
			(0.559)	(0.587)
Time Delay × Harshness			-0.549	-0.468
			(0.552)	(0.568)
Motivated Delay × Harshness			-0.573	-0.571
			(0.566)	(0.595)
Age		0.003		0.003
		(0.003)		(0.003)
Female		-0.011		-0.010
		(0.091)		(0.091)
Education (in years)		-0.039**		-0.038**
		(0.017)		(0.017)
Risk tolerance		0.303***		0.300***
		(0.049)		(0.049)
Positive Reciprocity		-0.104*		-0.102*
		(0.058)		(0.058)
Negative Reciprocity		0.059		0.058
		(0.050)		(0.050)
Altruism		0.136**		0.141***
		(0.054)		(0.054)
Trust		0.052		0.051
		(0.045)		(0.045)
Constant	-0.283**	0.124	-0.420**	-0.002
	(0.120)	(0.315)	(0.193)	(0.359)
Observations	1070	1055	1070	1055
Pseudo R^2	0.015	0.070	0.017	0.071

Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B2: OLS regressions examining beliefs (compliants only)

	(1)	(2)	(3)	(4)
Time Pressure	2.835	3.500	-0.206	-0.144
	(2.212)	(2.213)	(5.345)	(5.366)
Time Delay	0.204	1.216	-1.199	-0.702
	(2.006)	(2.011)	(4.962)	(4.993)
Motivated Delay	0.232	0.493	-0.932	-1.756
	(1.977)	(1.976)	(5.033)	(5.069)
Harshness	-7.735**	-7.365**	-10.644	-11.490
	(3.505)	(3.504)	(7.529)	(7.571)
Time Pressure × Harshness			6.663	7.992
			(10.610)	(10.674)
Time Delay × Harshness			3.037	4.171
			(9.909)	(9.944)
Motivated Delay × Harshness			2.521	4.866
			(10.040)	(10.026)
Age		-0.098*		-0.099*
		(0.057)		(0.057)
Female		-0.563		-0.558
		(1.532)		(1.534)
Education (in years)		-0.324		-0.332
		(0.275)		(0.275)
Risk tolerance		3.070***		3.081***
		(0.770)		(0.770)
Positive Reciprocity		-0.791		-0.814
		(0.995)		(0.996)
Negative Reciprocity		0.598		0.589
		(0.857)		(0.859)
Altruism		1.696*		1.673*
		(0.914)		(0.923)
Trust		-0.109		-0.067
		(0.777)		(0.783)
Constant	48.741***	56.668***	50.084***	58.728***
	(2.195)	(5.398)	(3.770)	(6.184)
Observations	1070	1055	1070	1055
R^2	0.007	0.036	0.007	0.037

Robust standard errors in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Additional tables and figures

Table B3: Probit regressions examining the effect of harshness on hostile behavior, divided by treatment

	Control		TP		TD		MD	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Harshness	-0.183	-0.179	0.063	-0.034	-0.732*	-0.635	-0.755*	-0.897**
	(0.382)	(0.409)	(0.370)	(0.383)	(0.401)	(0.423)	(0.403)	(0.440)
Age		0.004		-0.000		0.000		0.012
		(0.007)		(0.006)		(0.007)		(0.008)
Female		0.141		-0.186		0.057		-0.045
		(0.177)		(0.163)		(0.192)		(0.192)
Education (in years)		-0.061*		-0.006		-0.013		-0.024
		(0.034)		(0.029)		(0.036)		(0.029)
Risk tolerance		0.269***		0.303***		0.198**		0.433***
		(0.093)		(0.083)		(0.095)		(0.099)
Positive Reciprocity		-0.185*		0.144		-0.242**		-0.137
		(0.107)		(0.112)		(0.118)		(0.124)
Negative Reciprocity		0.099		-0.067		0.253**		-0.081
		(0.100)		(0.090)		(0.109)		(0.106)
Altruism		0.255**		-0.079		0.231*		0.319***
		(0.112)		(0.101)		(0.126)		(0.124)
Trust		0.117		-0.038		0.056		0.130
		(0.087)		(0.079)		(0.094)		(0.102)
Constant	-0.420**	0.138	-0.273	-0.071	-0.378*	-0.283	-0.427**	-0.535
	(0.192)	(0.642)	(0.187)	(0.540)	(0.197)	(0.666)	(0.198)	(0.565)
Observations	280	273	280	278	280	277	280	276

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table B4: Probit regressions examining the effect of harshness on hostile behavior, divided by delay

	No delay		Delay	
	(1)	(2)	(3)	(4)
Harshness	-0.054	-0.088	-0.743***	-0.709**
	(0.265)	(0.273)	(0.284)	(0.300)
Female		-0.049		0.002
		(0.118)		(0.133)
Age		0.003		0.005
		(0.005)		(0.005)
Education (in years)		-0.035		-0.021
		(0.022)		(0.022)
Risk tolerance		0.276***		0.309***
		(0.061)		(0.067)
Positive Reciprocity		-0.025		-0.188**
		(0.075)		(0.084)
Negative Reciprocity		0.012		0.099
		(0.065)		(0.073)
Altruism		0.056		0.253***
		(0.072)		(0.086)
Trust		0.016		0.062
		(0.057)		(0.066)
Constant	-0.346***‘	0.041	-0.403***	-0.313
	(0.134)	(0.402)	(0.140)	(0.419)
Observations	560	551	560	553

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure B1: Proportion of Hawk choices by treatment and harshness level

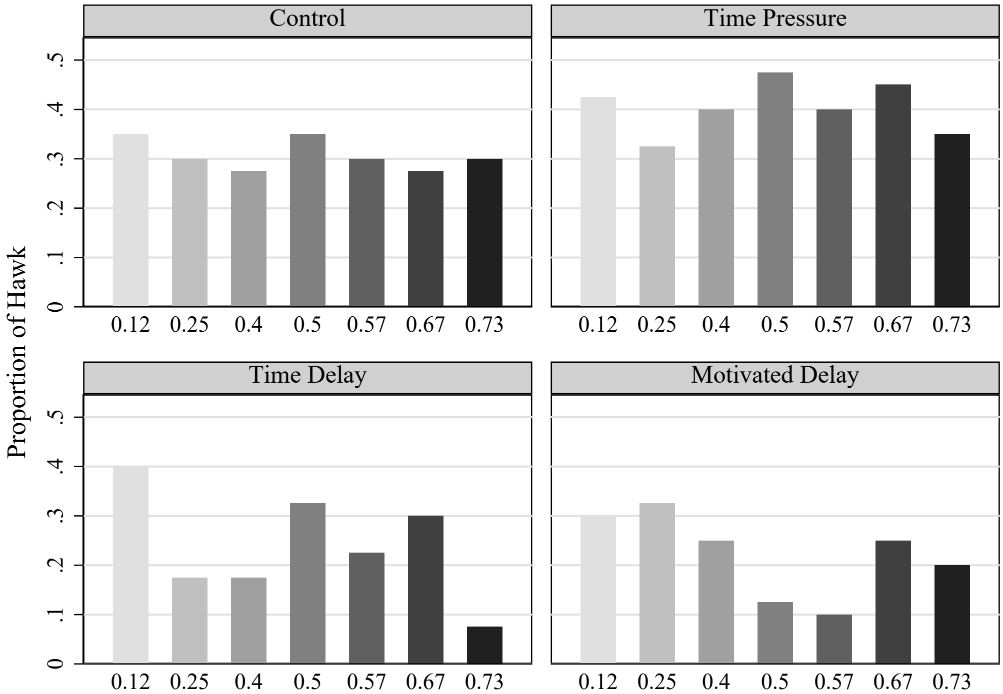

Figure B2: Proportion of Hawk choices by delayed vs. non-delayed subjects and harshness level

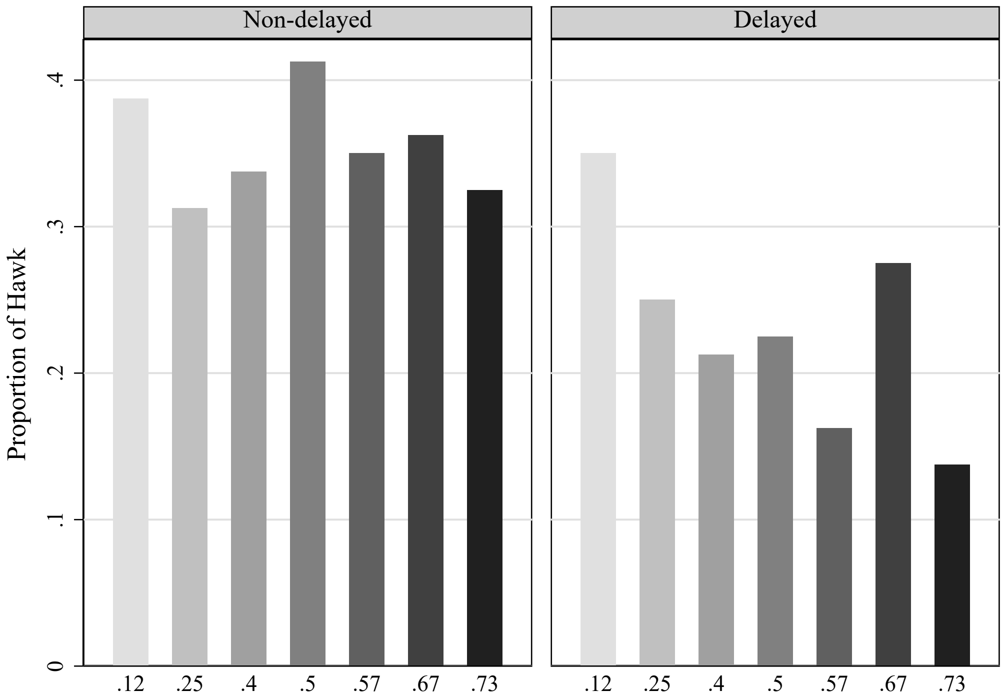

Figure B3: Proportion of best reply choices by treatment and harshness level

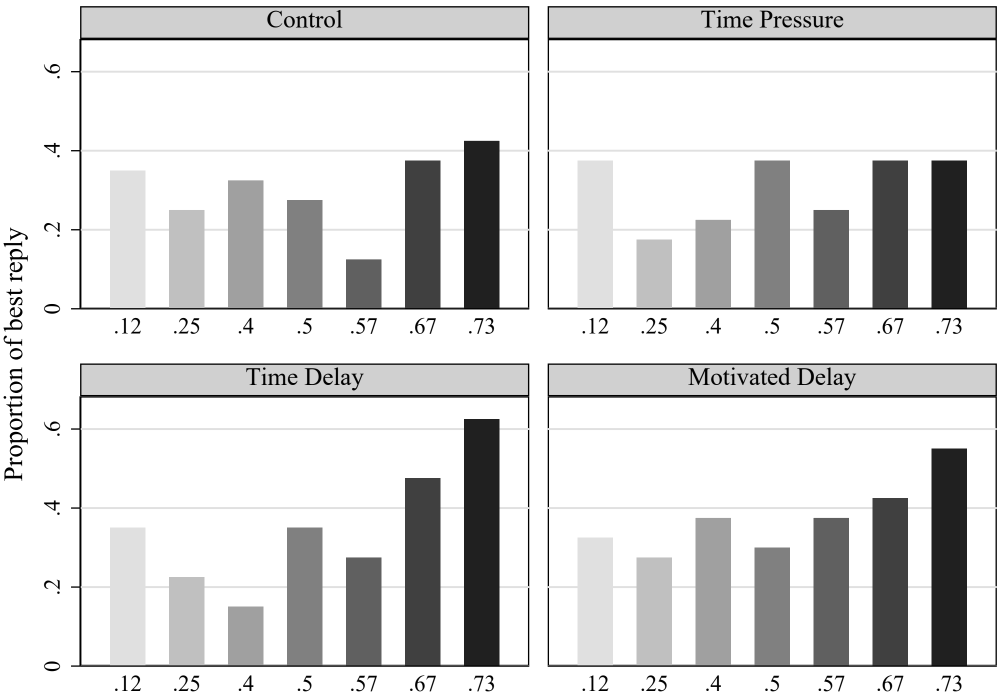

Figure B4: Proportion of best reply choices by delayed vs. non-delayed subjects and harshness level

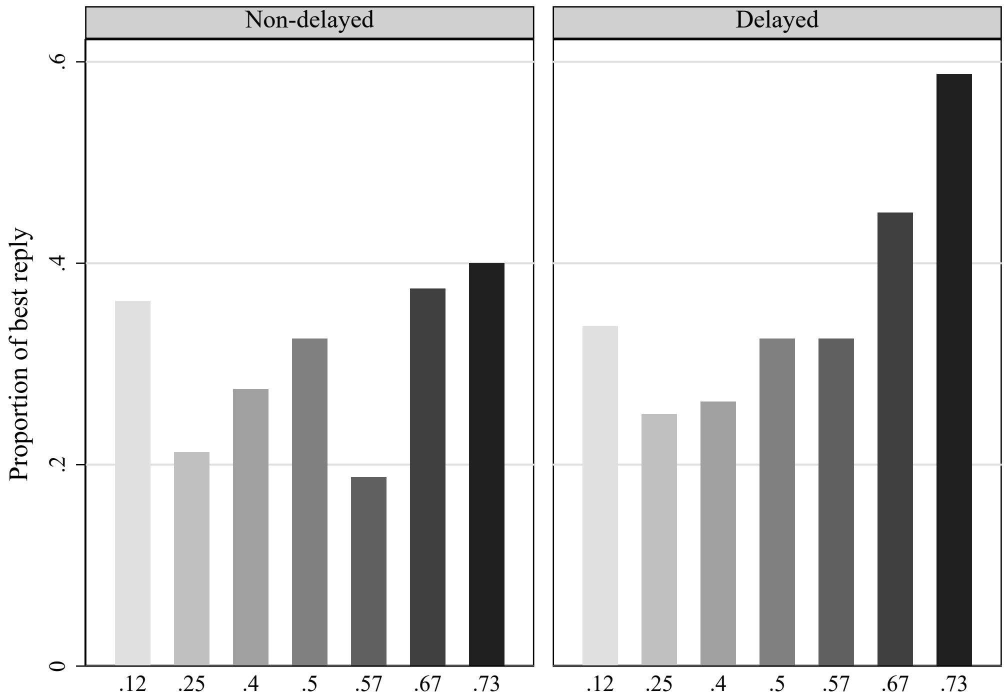

References

Achtziger, A., Alós-Ferrer, C., and Wagner, A. K. (2015). Money, depletion, and prosociality in the dictator game. *Journal of Neuroscience, Psychology, and Economics*, 8(1):1.

Achtziger, A., Alós-Ferrer, C., and Wagner, A. K. (2016). The impact of self-control depletion on social preferences in the ultimatum game. *Journal of Economic Psychology*, 53:1–16.

Arigapudi, S. and Heller, Y. (2025). Stable mixing in hawk–dove games under best experienced payoff dynamics. *Games and Economic Behavior*, 151:148–161.

Baumeister, R. F., Bratslavsky, E., Muraven, M., and Tice, D. M. (2018). Ego depletion: Is the active self a limited resource? In *Self-regulation and self-control*, pages 16–44. Routledge.

Bear, A. and Rand, D. G. (2016). Intuition, deliberation, and the evolution of cooperation. *Proceedings of the National Academy of Sciences*, 113(4):936–941.

Belloc, M., Bilancini, E., Boncinelli, L., and D'Alessandro, S. (2019). Intuition and deliberation in the stag hunt game. *Scientific reports*, 9(1):14833.

Benndorf, V. and Martinez-Martinez, I. (2017). Perturbed best response dynamics in a hawk–dove game. *Economics Letters*, 153:61–64.

Benndorf, V., Martinez-Martinez, I., and Normann, H.-T. (2016). Equilibrium selection with coupled populations in hawk–dove games: Theory and experiment in continuous time. *Journal of Economic Theory*, 165:472–486.

Bigoni, M., Blazquez De Paz, M., and Le Coq, C. (2026). Beyond hawks and doves: Can inequality ease coordination? *Economic Theory*.

Bilancini, E., Boncinelli, L., and Celadin, T. (2022a). Social value orientation and conditional cooperation in the online one-shot public goods game. *Journal of Economic Behavior & Organization*, 200:243–272.

Bilancini, E., Boncinelli, L., and Celadin, T. (2024). Manipulating response times in the cognitive reflection test: Time delay boosts deliberation, time pressure hinders it. *Journal of Behavioral and Experimental Economics*, 112:102273.

Bilancini, E., Boncinelli, L., and Di Paolo, R. (2025). Manipulating cognition in the one-shot stag-hunt game played online. *Unpublished results*.

Bilancini, E., Boncinelli, L., Guarnieri, P., and Spadoni, L. (2023). Delaying and motivating decisions in the (bully) dictator game. *Journal of Behavioral and Experimental Economics*, 107:102–106.

Bilancini, E., Boncinelli, L., Ille, S., and Vicario, E. (2022b). Memory retrieval and harshness of conflict in the hawk–dove game. *Economic Theory Bulletin*, 10(2):333–351.

Bilancini, E., Marcos-Prieto, P., Ordali, E., and Roni, R. (Unp.R.). Hostility is (almost) gender neutral. *Unpublished results.*

Blanco, M., Engelmann, D., Koch, A. K., and Normann, H.-T. (2010). Belief elicitation in experiments: is there a hedging problem? *Experimental economics*, 13:412–438.

Bouwmeester, S., Verkoeijen, P. P. J. L., Aczel, B., Barbosa, F., Bègue, L., Brañas-Garza, P., Chmura, T. G. H., Cornelissen, G., Døssing, F. S., Espín, A. M., Evans, A. M., Ferreira-Santos, F., Fiedler, S., Flegr, J., Ghaffari, M., Glöckner, A., Goeschl, T., Guo, L., Hauser, O. P., Hernan-Gonzalez, R., Herrero, A., Horne, Z., Houdek, P., Johannesson, M., Koppel, L., Kujal, P., Laine, T., Lohse, J., Martins, E. C., Mauro, C., Mischkowski, D., Mukherjee, S., Myrseth, K. O. R., Navarro-Martínez, D., Neal, T. M. S., Novakova, J., Pagà, R., Paiva, T. O., Palfi, B., Piovesan, M., Rahal, R.-M., Salomon, E., Srinivasan, N., Srivastava, A., Szaszi, B., Szollosi, A., Thor, K. Ø., Tinghög, G., Trueblood, J. S., Bavel, J. J. V., van 't Veer, A. E., Västfjäll, D., Warner, M., Wengström, E., Wills, J., and Wollbrant, C. E. (2017). Registered Replication Report: Rand, Greene, and Nowak (2012). *Perspectives on Psychological Science*, 12(3):527–542. PMID: 28475467.

Cappelletti, D., Güth, W., and Ploner, M. (2011). Being of two minds: Ultimatum offers under cognitive constraints. *Journal of Economic Psychology*, 32(6):940–950.

Capraro, V. (2017). Does the truth come naturally? Time pressure increases honesty in one-shot deception games. *Economics Letters*, 158:54–57.

Capraro, V. (2024). The dual-process approach to human sociality: Meta-analytic evidence for a theory of internalized heuristics for self-preservation. *Journal of Personality and Social Psychology*, 126(5):719.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Dohmen, T., Huffman, D., Schupp, J., Falk, A., Sunde, U., and Wagner, G. G. (2011). Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences. *Journal of the European Economic Association*, 9(3):522–550.

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2018). Global evidence on economic preferences. *The Quarterly Journal of Economics*, 133(4):1645–1692.

Falk, A., Becker, A., Dohmen, T., Huffman, D., and Sunde, U. (2023). The preference survey module: A validated instrument for measuring risk, time, and social preferences. *Management Science*, 69(4):1935–1950.

Feltovich, N. (2011). The effect of subtracting a constant from all payoffs in a hawk-dove game: Experimental evidence of loss aversion in strategic behavior. *Southern Economic Journal*, 77(4):837–850.

Gilbert, D. T., Pelham, B. W., and Krull, D. S. (1988). On cognitive busyness: When person perceivers meet persons perceived. *Journal of personality and social psychology*, 54(5):733.

Grimm, V. and Mengel, F. (2011). Let me sleep on it: Delay reduces rejection rates in ultimatum games. *Economics Letters*, 111(2):113–115.

Kocher, M. G., Pahlke, J., and Trautmann, S. T. (2013). Tempus fugit: Time pressure in risky decisions. *Management Science*, 59(10):2380–2391.

Kvarven, A., Strømland, E., Wollbrant, C., Andersson, D., Johannesson, M., Tinghög, G., Västfjäll, D., and Myrseth, K. O. R. (2020). The intuitive cooperation hypothesis revisited: A meta-analytic examination of effect size and between-study heterogeneity. *Journal of the Economic Science Association*, 6(1):26–42.

Liu, C.-J. and Hao, F. (2011). An application of a dual-process approach to decision making in social dilemmas. *The American Journal of Psychology*, 124(2):203–212.

Maynard Smith, J. (1982). *Evolution and the Theory of Games.* Cambridge University Press, Cambridge.

Merkel, A. L. and Lohse, J. (2019). Is fairness intuitive? an experiment accounting for subjective utility differences under time pressure. *Experimental Economics*, 22:24–50.

Oprea, R., Henwood, K., and Friedman, D. (2011). Separating the hawks from the doves: Evidence from continuous time laboratory games. *Journal of Economic Theory*, 146(6):2206–2225.

Ordali, E., Marcos-Prieto, P., Avvenuti, G., Ricciardi, E., Boncinelli, L., Pietrini, P., Bernardi, G., and Bilancini, E. (2024). Prolonged exertion of self-control causes increased sleep-like frontal brain activity and changes in aggressivity and punishment. *Proceedings of the National Academy of Sciences*, 121(47):e2404213121.

Palan, S. and Schitter, C. (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17:22–27.

Rand, D. G., Brescoll, V. L., Everett, J. A., Capraro, V., and Barcelo, H. (2016). Social heuristics and social roles: Intuition favors altruism for women but not for men. *Journal of Experimental Psychology: General*, 145(4):389.

Rand, D. G., Greene, J. D., and Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416):427–430.

Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., and Greene, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, 5(1):3677.

Saqib, N. U. and Chan, E. Y. (2015). Time pressure reverses risk preferences. *Organizational Behavior and Human Decision Processes*, 130:58–68.

Shenhav, A., Rand, D. G., and Greene, J. D. (2012). Divine intuition: Cognitive style influences belief in God. *Journal of Experimental Psychology: General*, 141(3):423.

Tinghög, G., Andersson, D., Bonn, C., Böttiger, H., Josephson, C., Lundgren, G., Västfjäll, D., Kirchler, M., and Johannesson, M. (2013). Intuition and cooperation reconsidered. *Nature*, 498(7452):E1–E2.

Wen, Y., Masschelein, S., and Ratan, A. (2022). Loss aversion in asymmetric anti-coordination games. *Southern Economic Journal*, 88(4):1549–1573.

Young, D. L., Goodie, A. S., Hall, D. B., and Wu, E. (2012). Decision making under time pressure, modeled in a prospect theory framework. *Organizational Behavior and Human Decision Processes*, 118(2):179–188.