

The Psychology of Negative-Sum Competition in Strategic Interactions

Christopher K. Hsee^{1, 2}, Ying Zeng³, Xilin Li⁴, and Alex Imas²

¹ Cheung Kong Graduate School of Business

² Booth School of Business, University of Chicago

³ Rotman School of Management, University of Toronto

⁴ China Europe International Business School

Many real-life examples—from interpersonal rivalries to international conflicts—suggest that people actively engage in competitive behavior even when it is negative sum (benefiting the self at a greater cost to others). This often leads to loss spirals where everyone—including the winner—ends up losing. Our research seeks to understand the psychology of such negative-sum competition in a controlled setting. To do so, we introduce an experimental paradigm in which paired participants have the option to repeatedly perform a behavior that causes a relatively small gain for the self and a larger loss to the other. Although they have the freedom not to engage in the behavior, most participants actively do so and incur substantial losses. We propose that an important reason behind the phenomena is shallow thinking—focusing on the immediate benefit to the self while overlooking the downstream consequences of how the behavior will influence their counterparts' actions. In support of the proposition, we find that participants are less likely to engage in negative-sum behavior, if they are advised to consider the downstream consequences of their actions, or if they are put in a less frenzied decision environment, which facilitates deeper thinking (acting in discrete vs. continuous time). We discuss how our results differ from prior findings and the implications of our research for mitigating negative-sum competition and loss spirals in real life.

Keywords: competition and cooperation, behavioral game theory, prisoner's dilemma, myopic decisions, System 1 versus System 2

Supplemental materials: <https://doi.org/10.1037/pspa0000344.supp>

Many important social interactions are characterized by competition over limited resources. Some competitive behavior is characterized as zero sum, generating positive outcomes for some parties and negative outcomes for others, where the two sides even out (Deutsch, 1949; Stanne et al., 1999). Competitive behavior can also be positive sum; namely, it creates more benefits than costs. In either of these cases, competitive behavior will not destroy overall value, and in the case of positive-sum behavior, it will actually increase overall value.

However, an important class of competitive behaviors is negative sum, where engaging in it leads to a larger net cost to one's counterpart than the benefit to oneself. Particularly, let x denote the payoff of an action to the self and y denote the payoff of the same action to one's

counterpart. We classify behavior as negative sum if its payoffs correspond to $x > 0 > y$ and $x + y < 0$. If both parties repeatedly engage in such behavior, both parties will incur more and more losses.

Life is replete with examples of people (or organizations) engaging in such negative-sum behaviors. Rival tribes engage in perpetuated blood feuds, hurting every tribe's livelihood. Rival companies repeatedly throw money into advertisement campaigns, reducing every company's net profits. Political candidates repeatedly slander each other, tarnishing everyone's reputation. Belligerent nations repeatedly attack each other, causing loss and suffering on every side.

While real-world examples offer suggestive evidence, they inevitably involve myriad confounding factors and can be explained in multiple ways. This research seeks to explore the psychology of negative-sum competitive behavior in a controlled setting. To do so, we develop a novel experimental paradigm that gives paired participants an opportunity to engage in a behavior that hurts the other more than it benefits the self. This paradigm allows us to examine whether people without prior animosity engage in the behavior and what factors motivate and moderate it. Our approach emulates extant research using simple "little-world" experiments to elucidate complex big-world problems (e.g., Asch, 1951; Axelrod & Dion, 1988; Shah et al., 2012).

Paradigm

We introduce an experimental paradigm that we refer to as the negative-sum game or simply NSG. In the NSG, two anonymous

This article was published Online First July 13, 2023.

Christopher K. Hsee  <https://orcid.org/0000-0002-8391-7847>

Ying Zeng  <https://orcid.org/0000-0001-9839-1182>

The authors thank Alyssa Eldrige, Bowen Ruan, Ryan Webb, and participants in George Wu's lab for their helpful suggestions at various stages of this project. Christopher K. Hsee was at Chicago Booth when the project started and at CKGSB when the project concluded.

All authors made important contributions in developing the research idea and designing the experiments. Christopher K. Hsee wrote the first draft of the paper, and all others revised the paper and approved the final version. Ying Zeng was instrumental in data collection.

Correspondence concerning this article should be addressed to Christopher K. Hsee, Booth School of Business, University of Chicago, 5807 South Woodlawn Avenue, Chicago, IL 60637, United States. Email: christopher.hsee@chicagobooth.edu

participants with no prior interactions are paired and go through the procedure remotely. In the beginning, each person is endowed with a sum of money (e.g., 100c). During the designated study period (e.g., 80s), each person has the option to perform a tedious behavior (e.g., type and enter the code “111222”) with a negative-sum effect; for example, every time the person enters the code, they gain 1c and the other participant loses 3c. Participants are told in advance that it is up to them whether and how many times to perform the behavior and that their earnings depend on their final balance.

As we will explain in greater detail later, the NSG is similar in payoff structure to the repeated prisoner’s dilemma. However, the NSG is more representative of real-world competitive environments in several ways: Unlike the repeated prisoner’s dilemma, which requires players to make one and only one defect-or-cooperate choice in each discrete round, the NSG allows people to make active, self-directed decisions at the time of their choosing.

Given the negative-sum nature of the behavior in the NSG, if both players perform the behavior and do so repeatedly, then everyone, including the supposed “winner,” will likely end up losing. Additionally, the group as a whole will lose more and more over time—a phenomenon we refer to as a loss spiral. On the other hand, if neither player performs the behavior, then both will walk away with their full endowment and achieve a superior outcome for both players.

When participants are given the freedom of whether or not to engage in the negative-sum behavior, will they engage in it? We predict that they will and that they will do so to the extent that everyone incurs significant losses.

We conjecture that an important driver of negative-sum behavior is *shallow thinking*—a tendency to focus on the immediate benefit to the self while overlooking longer term strategic consequences. More specifically, shallow thinking will lead people to act myopically in repeated social interactions, taking actions that maximize their own early rewards without thinking through the subsequent effects of those actions on the behavior of others. In the context of the NSG, such myopia would manifest as initiating and engaging in negative-sum behavior in order to secure an immediate higher payoff for the self, without considering that this behavior will lead one’s partner to respond accordingly.

We argue that shallow thinking is a fast, heuristic process (System 1) rather than a slow, deliberate one (System 2). In strategic environments like the NSG, shallow thinking predicts that people will overlook that engaging in the behavior may provoke one’s counterpart to retaliate; the resulting cycle of retaliation could eventually lead both to lose relative to their starting point. Since the immediate outcome of the negative-sum behavior is positive for the self while the negative consequences lay downstream of the initial decision, shallow thinking thus predicts that people will engage in the behavior.

The notion of shallow thinking builds on two existing lines of research. One is myopia, showing that decisionmakers focus on immediate outcomes and overlook long-term effects (Ameriks et al., 2007; Hofmann et al., 2012; Hsee et al., 2003; Kirby & Herrnstein, 1995; Laibson, 1997; Loewenstein & Prelec, 1992;

O’Donoghue & Rabin, 2015; Shafir et al., 1997). Myopia manifests itself in many ways. For example, people have been shown to select an immediate reward over a larger and delayed reward (e.g., prefer \$100 now over \$110 tomorrow), even though if a constant period is added to both options, they would prefer the latter (e.g., prefer \$110 in 11 days over \$100 in 10 days; Kirby & Herrnstein, 1995; Laibson, 1997; Loewenstein & Prelec, 1992). People also engage in immediately gratifying activities (e.g., smoking and eating unhealthy foods) while seemingly overlooking the long-term aversive consequences (Ameriks et al., 2007; Hofmann et al., 2012).

The idea of shallow thinking also draws on extant research showing egocentrism and a lack of perspective taking (e.g., Camerer et al., 1989; Crawford et al., 2008; Hsee et al., 2021; Madarász, 2012; Nagel, 1995; Ross et al., 1977). For example, people tend to consider their own values as more common and widespread than alternatives—a phenomenon termed the false-consensus effect (Krueger & Clement, 1994; Marks & Miller, 1987; Ross et al., 1977). Relatedly, past research shows that decisionmakers in strategic interactions tend to overlook the preference and beliefs of their counterparts, or mistakenly assume that their counterparts are less sophisticated than themselves (Arad & Rubinstein, 2012; Costa-Gomes & Crawford, 2006; Crawford et al., 2008; Hsee et al., 2021; Nagel, 1995).

Both myopia and egocentrism suggest that shallow thinking in strategic interactions will prompt people to engage in negative-sum behavior. People will focus on the immediate benefit to the self while overlooking the possible retaliatory behavior of their counterpart and its downstream negative consequences. Notably, shallow thinking can explain both why one would *initiate* the behavior (i.e., being the first in the NSG to enter the code) and why one would *repeat* the behavior (i.e., entering the code after their counterpart has done so). People are predicted to initiate the behavior because the immediate outcome of the behavior is beneficial to them. They would repeat the behavior for the same reason: The immediate outcome continues to be individually beneficial.

Besides shallow thinking, backward induction (e.g., Aumann, 1995; Dal Bó & Fréchette, 2019) is another potential explanation for why people would engage in negative-sum behavior in the NSG. It is in each player’s best interest to perform the behavior in the last moment of the game; using backward induction, it is thus also in their best interest to perform the behavior in every preceding moment, including the very first period.

Although backward induction and shallow thinking make similar predictions, they are psychologically very distinct. Standard game-theoretic approaches of backward induction require rationality, namely, that people are far-sighted and have accurate beliefs about the strategic responses of their counterparts. By contrast, shallow thinking presumes that people are boundedly rational—they are myopic in focusing on the immediate outcomes to the self while overlooking the strategic, long-term consequences of their actions.

Shallow thinking is a more likely driver of negative-sum behavior than backward induction. As outlined above, the idea

of shallow thinking builds on extensive empirical research on myopia and egocentrism; at the same time, there exists little empirical evidence that people (without training in game theory) engage in backward induction (e.g., Colman, 2003; Johnson et al., 2002; Levitt et al., 2011). Furthermore, as we report below, our finding that deeper thinking prompts reduce negative-sum behavior supports the shallow-thinking account; in contrast, if people behaved in line with backward induction, such a prompt would not affect behavior.

Another possible reason for engagement in negative-sum behavior is *retaliation*—to reciprocate what their counterpart has done to them (e.g., Nowak & Sigmund, 1993; Sigmund, 2007). Unlike backward induction, which is antithetical to shallow thinking, retaliation and shallow thinking are not mutually exclusive; the two motives likely coexist as contributors to negative-sum behavior. Notably, however, retaliation can only explain why people *repeat* the behavior; it cannot explain why they initiate the behavior in the first place. We will discuss other potential explanations in the General Discussion section.

We now proceed to outline testable predictions of the shallow-thinking account. If the engagement in negative-sum behavior is indeed driven at least in part by shallow thinking, then adding a piece of deeper thinking advice that prompts participants to consider the downstream consequences of their decisions in lieu of immediate payoffs would moderate negative-sum behavior. Shallow thinking also predicts that the level of potential “frenzy” in the decision environment will moderate negative-sum behavior.¹ We define a frenzied decision environment as one in which a person’s payoffs depend (or are perceived as depending) on how fast they act; in contrast, a nonfrenzied decision environment is one in which a person’s payoffs do not depend (or are perceived as not depending) on how fast they act. Importantly, whether a decision environment is frenzied or not is not a binary distinction but rather a continuum. By “frenzied” and “nonfrenzied,” we mean “relatively frenzied” and “relatively nonfrenzied”: The more one’s payoff depends on (or is perceived as depending on) how fast one acts, the more frenzied the decision environment.

If shallow thinking indeed plays an important role in engagement in negative-sum behavior, then there will be less engagement in negative-sum behavior as the decision environment becomes less prone to frenzy. This is because a frenzied environment prompts a sense of urgency, leading people to base their decision on System 1 processes and heuristics such as myopia and egocentrism (e.g., Zhu et al., 2018), while a nonfrenzied environment prompts deliberation and provides opportunities to mobilize System 2 processes such as perspective taking and cost-benefit analysis.

We argue that, relative to other paradigms used to study strategic interactions, the NSG is a frenzied decision environment. Recall that in the NSG, each player can incessantly perform the negative-sum behavior (e.g., enter the code 111222) at any speed they are able to, and their payoffs depend directly on the speed of their actions. This environment prompts a sense of urgency, which we argue leads to a greater reliance on fast heuristics and, thus, shallow thinking.

Shallow thinking predicts that making the NSG less frenzied will attenuate the prevalence of negative-sum behavior—even when the underlying payoff structure remains the same. Specifically, we compare the default NSG with a discretized NSG. Unlike the default version, in which players can incessantly perform the negative-sum

behavior, participants in the discretized NSG have to go through many discrete rounds where they can choose to engage in negative-sum behavior or not. In each round, they can take as long as they want to decide whether to engage in the behavior or not; after each choice, participants move on to the next round. The discretized NSG is a less frenzied decision environment than the default NSG because each player can enter the code only once in each round, and their payoff does not depend on how fast they do so. If the shallow-thinking account holds, people will be less likely to engage in negative-sum behavior in the discretized NSG than in the default NSG.

Because the underlying payoff structure of the default NSG is isomorphic to the discretized NSG, rational accounts such as backward induction predict similar levels of negative-sum behavior. In turn, if we indeed find higher rates of negative-sum behavior in the default NSG than in the discretized NSG, as we predicted, it provides evidence for shallow thinking as an important driver of the behavior.

Notably, the discretized NSG has the same basic structure as a repeated prisoner’s dilemma, in which defection (relative to cooperation) is negative sum. Thus, our above proposition predicts that people in our paradigm (the default NSG) are more likely to engage in negative-sum behavior than people in the repeated prisoner’s dilemma are likely to defect. Notably, this disparity would explain the difference between the current research and prior empirical work on repeated prisoner’s dilemmas, which has found fairly high levels of cooperation (see Embrey et al., 2018, for review).

Study Overview

We examined negative-sum behavior using the NSG and its variants. All studies were incentive compatible. Study 1 demonstrated a high prevalence of negative-sum behavior, higher than what one would expect from related findings in the literature on repeated prisoner’s dilemmas. Study 2 replicated Study 1’s results using different parameters and participants from a different population. Study 3 tested the shallow-thinking account by adding deeper thinking advice. Study 4 tested the shallow-thinking account by discretizing the NSG. Study 5 tested the combined effects of deeper thinking advice and discretization. In addition, we conducted two other studies to explore the moderators that we discuss in the General Discussion section and report the studies in [Supplemental Materials Section A](#). See [Table 1](#), for a summary of all studies and their key findings.

Regarding sample size, we preset a target sample size of 100 for studies with one condition and a target sample size of 200 for studies with two conditions. We increased the target sample size in Study 3 because we expected weaker effects of the manipulations, and we preregistered this. We advertised each study online or in the lab, asking interested participants to sign in during a certain time period and promising them a base payment plus possible additional earnings. After a participant signed in, we first asked them to go through a screening session to screen out bots and ensure that they had no technical difficulty (e.g., slow internet). During the screening session, participants were asked to enter a code multiple times (e.g., enter “111222” 10 times). After passing the screening task,

¹ The authors thank Clayton Critcher for suggesting the notion of “frenzy”.

Table 1
Overview of All Studies and Their Key Findings

Study	Objective(s)	Condition	Negative-sum behavior count	
			<i>Mdn</i>	<i>M (SD)</i>
Study 1	Tests for negative-sum behavior		45	42.17 (20.68)
Study 2	Tests replicability		37	34.48 (15.96)
Study 3	Tests the moderating effect of deeper thinking	Control	35	35.3 (19.36)
		Deeper thinking	29.5	30.03 (20.15)
Study 4	Tests the moderating effect of discretization	Default NSG	25	19.90 (9.76)
		Discretized NSG	6.5	9.96 (9.64)
Study 5	Tests the moderating effects of both deeper thinking and discretization	Default NSG/control	20	11.78 (9.75)
		Discretized NSG/control	6	8.44 (8.25)
		Default NSG/deeper thinking	3	9.44 (9.63)
		Discretized NSG/deeper thinking	2	7.17 (7.93)
Study S1	Tests the moderating effect of automation	Default NSG	58.5	57.86 (34.13)
		Automated NSG	44.5	44.67 (24.88)
Study S2	Tests the moderating effect of invisibility	Default NSG	33.5	29.53 (24.95)
		Invisible NSG	42	36.23 (23.35)

Note. NSG = negative-sum game.

participants began the study.² At the end of the study, participants reported demographic information such as gender and age; there were no systematic gender or age effects (see [Supplemental Materials Section B](#), for a summary of the results by gender).

Finally, we note that the data across all studies violated the normality assumption required for parametric analyses (Shapiro–Wilk normality test, all $ps < .05$). We, therefore, employed nonparametric Kernel regression as our statistical test throughout and reported medians in the main text accordingly. All data have been made available at the Open Science Framework and can be accessed with https://osf.io/vewf3/?view_only=a4a3ce81e3fd466c8e54fb07f0b289c0. All the materials are reported in the article.

Study 1: High Prevalence of Negative-Sum Behavior

Method

Study 1 tested whether people would engage in negative-sum behavior and fall into a loss spiral. The study was preregistered at <https://aspredicted.org/blind.php?x=4du98q>. Participants were 104 workers (47.1% female, 52.9% male; $M_{age} = 36.87$) recruited on MTurk in the United States. All participants received the following instructions:

In this study, you will be paired with another online worker to play a special game. You are anonymous to each other. For ease of reference, let us call the other participant Ze. You and Ze are given the same instructions.

At the beginning of the game, we give you 100 cents (\$1.00). You may gain or lose money during the game. After we finish collecting data, we will pay you the amount of money you have at the end of the game. This is real, not hypothetical.

The game will last at least 60 seconds, and it may last longer. The computer will tell you when the game ends.

Once the game begins, you have the option to enter the code 111222.

Every time you enter the code, you will gain 1 cent and Ze will lose 2 cents. The same is true for Ze.

It is up to you whether you enter the code and how many times you enter it.

If you want to enter the code, you must type it yourself; you must not use any shortcuts. If you don't want to enter the code, you may sit back and relax, but you should NOT exit the program before the game ends.

After reading the instructions, participants were matched into pairs. They then began the study, during which they saw the computer interface, as shown in [Figure 1](#). The statistics on the interface changed every time either player entered the code.

The game lasted for 100 s, though participants did not know the exact duration ahead of time. Upon completion, we paid every participant according to their final balance plus the advertised base payment.

Results and Discussion

The results show a high prevalence of negative-sum behavior and a steep loss spiral. On average, participants performed the behavior

² At the beginning, each participant was first paired anonymously with another participant. If a participant waited in vain for 3 min without being paired with another participant (because no one else signed in during that period), we gave the participant the option to leave with the base payment or continue to wait. Once two participants were paired, they began the study. The above procedure does not apply to either condition of Study 5 or the invisibility condition of [Supplemental Study S2](#), because those conditions did not involve simultaneous pairing of participants. If a participant dropped out before the study ended, we stopped the study and excluded both participants in the pair (if any). This led to the exclusions of 4, 0, 10, 4, 6, 16, and 10 participants in Studies 1–5, [Supplemental Study S1](#), and [S2](#), respectively. No one dropped out in Study 2 because the study was conducted in the lab. In each study, we continued recruiting new participants until the number of participants who had completed the study reached the target sample size. After we stopped recruiting new participants, we still allowed participants who had already signed in but had not completed the study to do so; therefore, the final sample size was somewhat larger than the target sample size in most studies. At the end of all studies, we collected participants' demographic information. Some studies also asked questions such as whether participants had experienced technical problems or had difficulty typing. Few participants said yes; we did not exclude any participants. Excluding participants who reported technical problems or typing difficulty does not meaningfully change the results. All studies were programed through LIONESS Lab ([Giamattei et al., 2020](#)).

Figure 1
Computer Interface During the NSG

The game has begun.

If you want to enter the code 111222, type it in the textbox and press Enter; otherwise, just sit back.

You have entered it 0 times.

Ze has entered it 0 times.

You now have 100 cents.

Ze now has 100 cents.

Note. NSG = negative-sum game. See the online article for the color version of this figure.

42.17 times (*Mdn*: 45 times). Consequently, they lost a significant amount (*M*: -\$42.17; *Mdn*: -\$39) even the “winner” in each pair (the one who ended up with relatively more money) lost a significant amount (*M*: -\$9.11; *Mdn*: -\$9.5, Wilcoxon $V = 374$, $p < .01$, compared with 0)—a Pyrrhic victory.

We also tracked and recorded participants’ balances every 10 s and presented the results in Figures 2 and 3. The figure reveals a steady loss spiral: both the relative winner and loser ended up losing more and more over time.


Notably, the prevalence of negative-sum behavior is higher than defection rates in the repeated prisoner’s dilemma as observed in previous research (e.g., Colman, 1995; Dal Bó & Fréchette, 2011; Embrey et al., 2018; Sally, 1995). We will discuss the difference between our paradigm and the repeated prisoner’s dilemma when reporting Study 4.

Study 2: A Replication

Method

Study 2 replicated Study 1 in a physical lab with university students as participants. It also included comprehension

Figure 2
Accumulated Typing Count Over Time in Study 1



Note. See the online article for the color version of this figure.

questions to ensure participants understood that it was up to them whether and how many times to perform the negative-sum behavior.

Participants were 102 students (64.7% female, 35.3% male; $M_{\text{age}} = 19.80$) recruited from a large public university in Canada, who completed the study in exchange for course credits. The study was conducted in a physical lab; every participant was seated in an individual cubicle with a computer in front of them. All participants received the following instructions:

You will be paired with another student to do this study. Both of you receive the following instructions:

The study lasts at least 70 seconds. For every 5 seconds after the 70th second, there is an even chance that the study will end or will continue.

Each of you have a virtual account. At the beginning of the study, you have \$100 in your account. During the study, you may gain or lose money depending on whether each of you enter the number 111222 and, if so, how many times each of you enter the number.

Specifically, every time one of you enter 111222, the person who enters it gains \$1 and the other person loses \$2.

It is entirely up to you whether you enter 111222 and how many times you enter it; we have no preference.

Your goal is to have as much money in your account as possible by the end of the study. After the entire project finishes, we will randomly select a pair of participants and actually pay them based on the money they have in their accounts at the end of the study. Therefore, the more money you have in your account at the end of the study, the more money you will actually receive if you are selected.

Please note:

1. Do not quit before the end of the study; otherwise, your data will be invalid.
2. You must not communicate with each other during the study.

Note that we never mentioned the word “game” or “player” in the instructions.


After reading the instructions, participants answered four comprehension questions, including one that tested their understanding that it was entirely up to them whether and how many times to enter the code (see Supplemental Materials Section C, for details). Each participant had to answer all comprehension questions correctly to proceed; if they answered any question incorrectly, they had to reanswer it until they were correct.

As noted in the instructions, the duration of the study was indefinite; we programed the study to last at least 70s, with a 50% chance of ending after every 5s interval after 70s. After the entire study finished, we randomly picked two participants and paid them based on their final balance, with 1 point equal to 1 cent (plus the base payment), as promised.

Results and Discussion

Again, the participants actively engaged in the negative-sum behavior: on average, participants performed the behavior 34.48 times (*Mdn*: 37 times). Consequently, the participants incurred significant losses (-\$34.48; *Mdn*: -\$34.5); even the “winner” lost

Figure 3
Accumulated Outcome Over Time in Study 1



Note. See the online article for the color version of this figure.

a significant amount ($Mdn: -\$21$, Wilcoxon $V = 169$, $p < .001$, compared with 0).

Study 2 replicated the results of Study 1. Since the study used a student sample, included a set of comprehension questions, and never mentioned the word “game” or “player,” it suggests that the results were not limited to online workers and not due to participants’ failure to understand our instructions or the use of specific words.

Study 3: The Moderating Effect of Deeper Thinking Advice

Method

In Study 3, we aimed to test our conjecture that shallow thinking drives negative-sum behavior in the NSG. To do so, we added a piece of deeper thinking advice to the paradigm, predicting that this would mitigate engagement in the behavior. The study was pre-registered at <https://aspredicted.org/blind.php?x=dv8mc9>. Participants were 308 workers (49.7% female, 50.3% male; $M_{age} = 40.58$) recruited on MTurk in the United States. Participants were randomly assigned to either a without-tip condition or a with-tip condition.

The procedure of the study was similar to that of Study 1, except for the following notable differences. The duration of the NSG was communicated to participants ahead of time (100s). Additionally, there was a cap on the number of times participants could repeat the negative-sum behavior (60 times). These changes make the payoff structure of NSG similar to that of the finitely repeated prisoner’s dilemma, where researchers have observed high rates of cooperation; observing active engagement in negative-sum behavior would therefore provide evidence for our shallow-thinking account as backward induction would operate similarly in the two contexts.

We endowed each participant with 1,000 points in the beginning and told them that every time they entered “111222,” they would gain 5 points and their counterpart (referred to as Ze) would lose 20 points, and that at the end of the study, we would randomly pick two participants and pay them a bonus according to the number of points they had at the end of the game, with 1 point worth 1 cent.

The study included two between-subjects conditions, one with deeper thinking advice and one without. The two conditions were identical, except that the with-tip condition showed participants the following advice before they entered the NSG:

Before the game begins, please take a moment to carefully think about what you will do during the game. When making your decision, don’t just think about its immediate outcome; you should think about its downstream consequences, that is, think about how your decision may influence Ze’s decision, how Ze’s decision may further influence your decision, and so on, and what the consequences of these decisions are.

In addition, the with-advice condition included a brief reminder during the game, which read, “When making your decision, don’t just think about its immediate outcome; you should think about its downstream consequences.”

After reading the instructions and before starting the study, participants completed a comprehension check about the instructions (see [Supplemental Materials Section C](#)). At the conclusion of the entire study, we randomly picked two participants and paid them based on their final balance, with 1 point equal to 1 cent, plus the base payment.

Results and Discussion

As predicted, the provision of the deeper thinking advice significantly reduced engagement in the negative-sum behavior: compared with participants in the without-advice condition, participants in the with-advice condition performed the behavior fewer times ($M: 30.03$ times vs. 35.30 times; $Mdn: 29.5$ times vs. 35.0 times, $z = -2.37$, $p = 0.02$) and lost fewer points ($M: -450.50$ vs. -529.56 points; $Mdn: -392.5$ vs. -535.0 points, $z = 1.97$, $p = .05$).^{3,4}

The results of this study support the shallow-thinking account and cast doubt on backward induction as a viable alternative process. Recall that the NSG in this study had a definite duration and a fixed cap for how many times one could type. Furthermore, since backward induction is presumably a deliberative (deeper thinking) process that predicts defection (i.e., typing), we should have observed *more* negative-sum behavior in the with-advice condition than in the without-advice condition—the opposite of what we found.

Study 4: The Moderating Effect of Discretizing the NSG

Method

Study 4 tested our prediction that people are less likely to engage in negative-sum behavior in the discretized NSG than in the default NSG. Participants were 216 workers (52.3% female, 47.7% male; $M_{age} = 38.43$) recruited on MTurk in the United States, and they were assigned to either a default NSG condition or a discretized NSG condition. We tried our best to make the two conditions as similar as possible. All participants received the following opening instructions:

³ To address the potential interdependence of data points within each pair, we also analyzed the data using hierarchical linear modeling, controlling for pair-level random intercepts, and we found similar results. See [Supplemental Materials Section D](#), for details.

⁴ Results in this study and those that follow come from nonparametric Kernel regressions. This regression is analogous to parametric regressions (e.g., linear regression) in modeling the mean of the outcome conditional on independent variables, but unlike parametric regressions, Kernel regression makes no assumptions about the functional form of this relationship. Our theory cannot predict the normality of data distributions in advance, and our preregistered analyses were parametric (e.g., analysis of variance). Results are similar in magnitude when conducting parametric tests.

In this study, you will be paired with another online worker to play a special game. You are anonymous to each other. For ease of reference, let's call the other participant Ze. You and Ze are receiving the same instructions, and all the rules that apply to you also apply to Ze.

At the beginning of the game, we give you 1,000 points. You may gain or lose points during the game. After we complete the entire project (which may take a few days), we will randomly select two participants and give them a bonus based on the points they have at the end of the game, with 1 point worth 1 cent. Therefore, the more points you have at the end of the game, the more money you will get if you are selected.

Participants in the default NSG condition then read:

The game lasts 100 seconds. During the game, you have the option to enter the code 111222. You may enter the code up to 25 times. Every time you enter the code, you will gain 2 points and Ze will lose 20 points. You don't have to enter the code. It is up to you whether you enter the code and how many times you enter it (up to 25 times).

Participants in the discretized NSG condition read the following instead:

The game consists of 25 rounds; each round lasts 4 seconds. During each round, you have the option to enter the code 111222. In any given round, you may enter the code only once. If you enter the code, you will gain 2 points and Ze will lose 20 points.

You don't have to enter the code. It is up to you whether you enter the code.

Finally, all participants read:

If you don't want to enter the code, you may sit back and relax, but you should NOT exit the program before the game is over.

Note that the two conditions had the same duration (100 s), the same type of negative-sum behavior (typing the code 111222), the same payoff structure (causing the self to gain 2 points and the other to lose 20 points), and the same upper limit of how many times one could type (25 times). The only difference was that in the default NSG condition, participants could incessantly type the code at any speed, but in the discretized NSG condition, participants had to go through 25 discrete rounds where they could type the code only once.

After reading the instructions, participants answered a set of comprehension questions (see [Supplemental Materials Section C](#)). Upon passing all the comprehension questions, they began the study. At the end of the entire study, we randomly picked two participants and paid them based on their final balance, with 1 point equaling 1 cent, plus the base payment.

Results and Discussion

As predicted, relative to participants in the default NSG condition, participants in the discretized NSG condition performed the negative-sum behavior fewer times (M : 9.96 times vs. 19.90 times; Mdn : 6.5 times vs. 25.0 times, $z = -7.59$, $p < .001$), and lost fewer points (M : -179.33 vs. -358.17 points; Mdn : -117.0 vs. -450.0 points, $z = -7.14$, $p < .001$).

Since the discretized NSG is analogous to a repeated prisoner's dilemma, the results of Study 4 may help explain why the defection

rates prior research found in the repeated prisoner's dilemma (e.g., [Embrey et al., 2018](#)) were usually lower than the prevalence of negative-sum behavior in our default NSG.

Together, the findings of Studies 3 and 4 cast doubt on backward induction as an explanation for engagement in negative-sum behavior. Backward induction would predict more engagement in negative-sum behavior in the discretized condition and the condition with deeper thinking advice since these manipulations should prompt players to think more "rationally" (i.e., strategically). Instead, we found the opposite results, which are consistent with the shallow-thinking account.

Study 5: The Combined Effects of Deeper Thinking Advice and Discretization

Method

Study 5 replicated both Studies 3 and 4 in the same setting by manipulating both deeper thinking and discretization in a 2 (deeper thinking advice: without vs. with) $\times$ 2 (type of NSG: default vs. discretized) between-subjects design. To minimize retaliation as a motive for engaging in negative-sum behavior, we designed the paradigm so that the players could not see the actions of their counterparts during the game, but their counterparts could.

We recruited 448 participants (46.2% female, 52.7% male, 1.1% other; $M_{age} = 38.91$) on CloudResearch. They were randomly assigned to one of the four conditions. All participants received the following opening instructions:

In this study, you will be paired with another online worker to play a special game. You are anonymous to each other. Unless otherwise specified, all the rules that apply to you also apply to the other player.

At the beginning of the game, we give you 1,000 points. You may gain or lose points during the game. After we complete the entire project (which may take a few days), we will randomly select two participants and give them a bonus based on the points they have at the end of the game, with 1 point worth 1 cent. Therefore, the more points you have at the end of the game, the more money you will get if you are selected.

Participants in the default NSG condition then read:

The game lasts 60 seconds. During the game, you may either type and enter the code 220, or just sit back.

You may enter the code 220 up to only 20 times during the game.

Every time you enter 220, you will gain 2 points and the other player will lose 20 points.

It is up to you whether and how many times you enter the code (up to 20 times).

During the game, you cannot see how many times the other player has entered the code, but the other player can see how many times you have entered the code.

Participants in the discretized NSG condition read the following instead:

The game consists of 20 rounds, and each round will last 3 seconds. In each round, you may either type and enter the code 220, or just sit back.

You may enter the code 220 up to only once in each round.

If you enter 220, you will gain 2 points and the other player will lose 20 points.

It is up to you whether you enter the code (up to once per round).

During the game, you cannot see how many times the other player has chosen each option, but the other player can see how many times you have chosen each option.

To make sure participants understood the instructions, we asked a set of comprehension questions (see [Supplemental Materials Section C](#)).

After answering all comprehension questions correctly, the participants proceeded to start the study. In the with-advice condition, participants in both the default and the discretized NSG conditions saw the following on their screen:

Tip: When you decide whether to enter the code 220 during the game, you should carefully think about the downstream consequences of your decision—e.g., how your decision will influence the other player's decision, and how their decision will influence the outcome of the game.

Note that participants in both conditions could perform the negative-sum behavior the maximum number of times if they wanted to because the maximum number of times in both conditions was only 20, the code was easy to type (220), and participants had enough time to type the code for the maximum number of times.

Upon finishing the data collection, we randomly selected two participants and paid each of them the maximum amount of the bonus that they could have earned based on their typing count.

Results and Discussion

Table 1 presents the results. Our analysis found a significant main effect of type of NSG ($z = 3.14, p = .002$), a marginally significant main effect of deeper thinking advice ($z = 1.70, p = .089$), and, critically, a significant interaction between type of NSG and deeper thinking advice ($z = 2.96, p = .003$).

Specifically, when no advice was given, participants in the default NSG again actively engaged in the negative-sum behavior ($M: 11.78$ times; $Mdn: 20$ times). Note that their median typing count (20) was the maximum number of times they could perform the negative-sum behavior. Since the players could not see the actions of their counterparts during the game, this result could not easily be attributed to retaliation. On the other hand, even without the deeper thinking advice, participants in the discretized-NSG condition performed the behavior significantly fewer times than participants in the default NSG condition ($M: 8.44$ times vs. 11.78 times; $Mdn: 6$ times vs. 20 times, $z = -2.67, p = .008$). This result replicated the finding of Study 4.

However, the difference between the default and the discretized NSG conditions disappeared when participants were given the deeper thinking advice ($M: 9.44$ times vs. 7.17 times; $Mdn: 3$ times vs. 2 times, $z = 1.07, p = .28$). These results indicate that prompting people to make more deliberative decisions leads to similar behavior in the default and discretized NSG. This provides further evidence for shallow thinking as the driver of negative-sum behavior in the default NSG.

General Discussion

Observations suggest that many human calamities—from interpersonal feuds to international wars—occur because people engage in negative-sum competitive behavior that is beneficial to the actor in the short run but detrimental to everyone in the long run. Despite the observations, there is little rigorous scientific work on the underlying drivers of negative-sum behavior. Our work seeks to shed light on the psychology of negative-sum behavior by introducing a novel experimental paradigm. Dubbed the NSG, this paradigm captures the frenzied nature of many real-world competitive situations that lead to loss spirals. This feature makes the NSG more ecologically valid than functionally similar classic games such as the repeated prisoner's dilemma. Using the NSG paradigm, we find that people actively engage in negative-sum behavior—despite knowing they do not have to—and incur significant losses as a consequence. We propose that the high prevalence of negative-sum behavior is driven at least in part by shallow thinking and identify two moderators that support our proposition—prompts that spur deeper thinking and discretization.

Besides shallow thinking, we considered two rational channels as possible alternative mechanisms: backward induction and retaliation. Our Studies 3–5 point against these mechanisms. The deeper thinking prompt and discretization should, if anything, increase negative-sum behavior if it was due to backward induction, as both manipulations should amplify rational strategic responses. Retaliation is a viable explanation for why people repeat negative-sum behavior in studies where participants could see the actions of their counterparts, but cannot explain why participants initiated the behavior in the first place, nor why participants repeated the behavior in Study 5 when they could not see the actions of their counterparts. Similarly, retaliation cannot explain the moderating effects.

We now consider three additional possible reasons. The first is overconfidence (Koriat et al., 1980; McKenzie, 1997; Soll, 1996). In the context of our paradigm, overconfidence refers to the belief that “I can type so much faster than my counterpart that I am not only able to beat my counterpart but also able to make an absolute gain at the end.” Another explanation is idleness: participants may engage in negative-sum behavior because they are otherwise idle and bored (e.g., Hsee et al., 2010; Pfattheicher et al., 2021; Westgate & Wilson, 2018). A third possible explanation is curiosity: Participants are curious about their counterpart's reactions (Hsee & Ruan, 2016; Loewenstein, 1994; Ruan et al., 2018). While these motives could explain why participants engage in the negative-sum behavior, none of them could explain any of the moderating effects.

Besides the deeper thinking prompt and discretization, other factors may also moderate negative-sum behavior. Study S1, reported in [Supplemental Materials Section A](#), investigated one such moderator. The study compared the default NSG with an automated version of the game in which players only needed to turn an OFF/ON switch to the ON position, and the computer would automatically and continuously produce the negative-sum outcome at a constant speed for them; if they wanted to stop the process, they would only need to turn the switch back to the OFF position. We predicted a lower prevalence of negative-sum behavior in the automated NSG than in the default NSG because payoffs in the automated NSG do not depend on how fast participants executed the negative-sum behavior (i.e., how fast they type the code). As such, the automated NSG would not create the same degree of urgency or

frenzy as the default NSG, and in turn, people would be less likely to engage in negative-sum behavior in the former than in the latter. This is indeed what we found in Study S1.

Note that the automated NSG is essentially a continuous prisoner's dilemma that has been studied extensively in the game-theory literature. The results of Study S1 explain why negative-sum behavior in the default NSG is more prevalent than defection in the continuous prisoner's dilemma, as documented in the existing literature (Friedman & Oprea, 2012). We believe that our paradigm (the default NSG) is more representative of most real-world competitive environments than the continuous prisoner's dilemma (the automated NSG) because people in real life typically need to exert effort actively to compete.

Another potential moderator is whether or not participants can see each other's actions during the interaction. In all our studies, participants were either able to see each other's actions (Studies 1–4) or told that their counterparts could see their actions (Study 5). What would happen if participants could not see each other's actions at all? We suspect that mutual invisibility may increase one's tendency to perform negative-sum behavior because she no longer needs to worry about retaliation from her counterpart. We conducted a study that manipulated whether participants could see each other's actions and found that invisibility indeed increased engagement in the negative-sum behavior; see Study S2 in Supplemental Materials Section A, for details.

Besides the specific moderating factors discussed above, the general social milieu (culture) may also influence people's tendency to engage in negative-sum behavior. We surmise that people are more likely to engage in negative-sum behavior and will lose more as a result in cultures that value competition and individual success than in cultures that value cooperation and joint benefit. This conjecture echoes the finding that people in the prisoner's dilemma are more likely to defect when the dilemma is labeled as a Wall Street Game than as a Community Game (Liberman et al., 2004).

Finally, individual differences may also influence the tendency to engage in negative-sum behavior. It would be interesting for future research to explore whether negative-sum behavior is related to people's degree of cognitive reflection (Frederick, 2005), intertemporal discount rate (Frederick et al., 2003), propensity for perspective taking (Epley et al., 2006), level of cognitive hierarchy (Camerer et al., 2004), social value orientation (Messick & McClintock, 1968; Van Lange & Kuhlman, 1994), or tendency to use feelings versus reason to guide decisions (Hsee et al., 2015).

Many competitive behaviors are positive sum, and engaging in positive-sum competitive behaviors generates positive values and likely benefits everyone involved. However, not all competitive behaviors are positive sum; some are zero-sum, and some are even negative sum (Dawes et al., 1977; Hsee et al., 2012; Rapoport & Chammah, 1965). This research highlights the dark side of competition and shows in a controlled experimental paradigm that people would actively engage in competitive behavior even if it is clearly negative sum and suffer significant mutual losses as a consequence. We hope that our work will stimulate further research on this topic and help reduce negative-sum outcomes in the real world.

References

Ameriks, J., Caplin, A., Leahy, J., & Tyler, T. (2007). Measuring self-control problems. *The American Economic Review*, 97(3), 966–972. <https://doi.org/10.1257/aer.97.3.966>

- Arad, A., & Rubinstein, A. (2012). The 11–20 money request game: A level-k reasoning study. *American Economic Review*, 102(7), 3561–3573. <https://doi.org/10.1257/aer.102.7.3561>
- Asch, S. E. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men; research in human relations* (pp. 177–190). Carnegie Press.
- Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior*, 8(1), 6–19. [https://doi.org/10.1016/S0899-8256\(05\)80015-6](https://doi.org/10.1016/S0899-8256(05)80015-6)
- Axelrod, R., & Dion, D. (1988). The further evolution of cooperation. *Science*, 242(4884), 1385–1390. <https://doi.org/10.1126/science.242.4884.1385>
- Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3), 861–898. <https://doi.org/10.1162/0033555041502225>
- Camerer, C. F., Loewenstein, G., & Weber, M. (1989). The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy*, 97(5), 1232–1254. <https://doi.org/10.1086/261651>
- Colman, A. M. (1995). Prisoner's Dilemma, Chicken, and mixed-strategy evolutionary equilibria. *Behavioral and Brain Sciences*, 18(3), 550–551. <https://doi.org/10.1017/S0140525X00039704>
- Colman, A. M. (2003). Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioral and Brain Sciences*, 26(2), 139–153. <https://doi.org/10.1017/S0140525X03000050>
- Costa-Gomes, M. A., & Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5), 1737–1768. <https://doi.org/10.1257/aer.96.5.1737>
- Crawford, V. P., Gneezy, U., & Rottenstreich, Y. (2008). The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *American Economic Review*, 98(4), 1443–1458. <https://doi.org/10.1257/aer.98.4.1443>
- Dal Bó, P., & Fréchet, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1), 411–429. <https://doi.org/10.1257/aer.101.1.411>
- Dal Bó, P., & Fréchet, G. R. (2019). Strategy choice in the infinitely repeated prisoner's dilemma. *The American Economic Review*, 109(11), 3929–3952. <https://doi.org/10.1257/aer.20181480>
- Dawes, R. M., McTavish, J., & Shaklee, H. (1977). Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and Social Psychology*, 35(1), 1–11. <https://doi.org/10.1037/0022-3514.35.1.1>
- Deutsch, M. (1949). A theory of cooperation and competition. *Human Relations*, 2(2), 129–152. <https://doi.org/10.1177/001872674900200204>
- Embrey, M., Fréchet, G. R., & Yuksel, S. (2018). Cooperation in the finitely repeated prisoner's dilemma. *The Quarterly Journal of Economics*, 133(1), 509–551. <https://doi.org/10.1093/qje/qjx033>
- Epley, N., Caruso, E., & Bazerman, M. H. (2006). When perspective taking increases taking: Reactive egoism in social interaction. *Journal of Personality and Social Psychology*, 91(5), 872–889. <https://doi.org/10.1037/0022-3514.91.5.872>
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19(4), 25–42. <https://doi.org/10.1257/089533005775196732>
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2003). Time discounting and time preference: A critical review. In G. Loewenstein, D. Read, & R. Baumeister (Eds.), *Time and decision: Economic and psychological perspectives on intertemporal choice* (pp. 13–86). Sage Publications.
- Friedman, D., & Oprea, R. (2012). A continuous dilemma. *The American Economic Review*, 102(1), 337–363. <https://doi.org/10.1257/aer.102.1.337>
- Giamattei, M., Yahosseini, K. S., Gächter, S., & Molleman, L. (2020). LIONESS Lab: A free web-based platform for conducting interactive experiments online. *Journal of the Economic Science Association*, 6(1), 95–111. <https://doi.org/10.1007/s40881-020-00087-0>

- Hofmann, W., Vohs, K. D., & Baumeister, R. F. (2012). What people desire, feel conflicted about, and try to resist in everyday life. *Psychological Science*, 23(6), 582–588. <https://doi.org/10.1177/0956797612437426>
- Hsee, C. K., & Ruan, B. (2016). The Pandora effect: The power and peril of curiosity. *Psychological Science*, 27(5), 659–666. <https://doi.org/10.1177/0956797616631733>
- Hsee, C. K., Shen, L., Zhang, S., Chen, J., & Zhang, L. (2012). Fate or fight: Exploring the hedonic costs of competition. *Organizational Behavior and Human Decision Processes*, 119(2), 177–186. <https://doi.org/10.1016/j.obhdp.2012.07.005>
- Hsee, C. K., Yang, A. X., & Wang, L. (2010). Idleness aversion and the need for justifiable busyness. *Psychological Science*, 21(7), 926–930. <https://doi.org/10.1177/0956797610374738>
- Hsee, C. K., Yang, Y., Zheng, X., & Wang, H. (2015). Lay rationalism: Individual differences in using reason versus feelings to guide decisions. *Journal of Marketing Research*, 52(1), 134–146. <https://doi.org/10.1509/jmr.13.0532>
- Hsee, C. K., Yu, F., Zhang, J., & Zhang, Y. (2003). Medium maximization. *The Journal of Consumer Research*, 30(1), 1–14. <https://doi.org/10.1086/374702>
- Hsee, C. K., Zeng, Y., Li, X., & Imas, A. (2021). Bounded rationality in strategic decisions: Undershooting in a resource pool-choice dilemma. *Management Science*, 67(10), 6553–6567. <https://doi.org/10.1287/mnsc.2020.3814>
- Johnson, E. J., Camerer, C., Sen, S., & Rymon, T. (2002). Detecting failures of backward induction: Monitoring information search in sequential bargaining. *Journal of Economic Theory*, 104(1), 16–47. <https://doi.org/10.1006/jeth.2001.2850>
- Kirby, K. N., & Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science*, 6(2), 83–89. <https://doi.org/10.1111/j.1467-9280.1995.tb00311.x>
- Koriat, A., Lichtenstein, S., & Fischhoff, B. (1980). Reasons for confidence. *Journal of Experimental Psychology: Human Learning and Memory*, 6(2), 107–118. <https://doi.org/10.1037/0278-7393.6.2.107>
- Krueger, J., & Clement, R. W. (1994). The truly false consensus effect: An ineradicable and egocentric bias in social perception. *Journal of Personality and Social Psychology*, 67(4), 596–610. <https://doi.org/10.1037/0022-3514.67.4.596>
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2), 443–478. <https://doi.org/10.1162/003355397555253>
- Levitt, S. D., List, J. A., & Sadoff, S. E. (2011). Checkmate: Exploring backward induction among chess players. *The American Economic Review*, 101(2), 975–990. <https://doi.org/10.1257/aer.101.2.975>
- Liberman, V., Samuels, S. M., & Ross, L. (2004). The name of the game: Predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. *Personality and Social Psychology Bulletin*, 30(9), 1175–1185. <https://doi.org/10.1177/0146167204264004>
- Loewenstein, G. F. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, 116(1), 75–98. <https://doi.org/10.1037/0033-2909.116.1.75>
- Loewenstein, G. F., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. *The Quarterly Journal of Economics*, 107(2), 573–597. <https://doi.org/10.2307/2118482>
- Madarász, K. (2012). Information projection: Model and applications. *The Review of Economic Studies*, 79(3), 961–985. <https://www.jstor.org/stable/23261376>
- Marks, G., & Miller, N. (1987). Ten years of research on the false-consensus effect: An empirical and theoretical review. *Psychological Bulletin*, 102(1), 72–90. <https://doi.org/10.1037/0033-2909.102.1.72>
- McKenzie, C. R. (1997). Underweighting alternatives and overconfidence. *Organizational Behavior and Human Decision Processes*, 71(2), 141–160. <https://doi.org/10.1006/obhd.1997.2716>
- Messick, D. M., & McClintock, C. G. (1968). Motivational bases of choice in experimental games. *Journal of Experimental Social Psychology*, 4(1), 1–25. [https://doi.org/10.1016/0022-1031\(68\)90046-2](https://doi.org/10.1016/0022-1031(68)90046-2)
- Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5), 1313–1326. <https://www.jstor.org/stable/2950991>
- Nowak, M., & Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner's Dilemma game. *Nature*, 364(6432), 56–58. <https://doi.org/10.1038/364056a0>
- O'Donoghue, T., & Rabin, M. (2015). Present bias: Lessons learned and to be learned. *American Economic Review*, 105(5), 273–279. <https://doi.org/10.1257/aer.p.20151085>
- Pfafftheicher, S., Lazarević, L. B., Westgate, E. C., & Schindler, S. (2021). On the relation of boredom and sadistic aggression. *Journal of Personality and Social Psychology*, 121(3), 573–600. <https://doi.org/10.1037/pspi0000335>
- Rapoport, A., & Chammah, A. M. (1965). Sex differences in factors contributing to the level of cooperation in the prisoner's dilemma game. *Journal of Personality and Social Psychology*, 2(6), 831–838. <https://doi.org/10.1037/h0022678>
- Ross, L., Greene, D., & House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13(3), 279–301. [https://doi.org/10.1016/0022-1031\(77\)90049-X](https://doi.org/10.1016/0022-1031(77)90049-X)
- Ruan, B., Hsee, C. K., & Lu, Z. Y. (2018). The teasing effect: An underappreciated benefit of creating and resolving an uncertainty. *JMR, Journal of Marketing Research*, 55(4), 556–570. <https://doi.org/10.1509/jmr.15.0346>
- Sally, D. (1995). Conversation and cooperation in social dilemmas: A meta-analysis of experiments from 1958 to 1992. *Rationality and Society*, 7(1), 58–92. <https://doi.org/10.1177/1043463195007001004>
- Shafir, E., Diamond, P., & Tversky, A. (1997). Money Illusion. *The Quarterly Journal of Economics*, 112(2), 341–374. <https://doi.org/10.1162/003355397555208>
- Shah, A. K., Mullainathan, S., & Shafir, E. (2012). Some consequences of having too little. *Science*, 338(6107), 682–685. <https://doi.org/10.1126/science.1222426>
- Sigmund, K. (2007). Punish or perish? Retaliation and collaboration among humans. *Trends in Ecology & Evolution*, 22(11), 593–600. <https://doi.org/10.1016/j.tree.2007.06.012>
- Soll, J. B. (1996). Determinants of overconfidence and miscalibration: The roles of random error and ecological structure. *Organizational Behavior and Human Decision Processes*, 65(2), 117–137. <https://doi.org/10.1006/obhd.1996.0011>
- Stanne, M. B., Johnson, D. W., & Johnson, R. T. (1999). Does competition enhance or inhibit motor performance: A meta-analysis. *Psychological Bulletin*, 125(1), 133–154. <https://doi.org/10.1037/0033-2909.125.1.133>
- Van Lange, P. A., & Kuhlman, D. M. (1994). Social value orientations and impressions of partner's honesty and intelligence: A test of the might versus morality effect. *Journal of Personality and Social Psychology*, 67(1), 126–141. <https://doi.org/10.1037/0022-3514.67.1.126>
- Westgate, E. C., & Wilson, T. D. (2018). Boring thoughts and bored minds: The MAC model of boredom and cognitive engagement. *Psychological Review*, 125(5), 689–713. <https://doi.org/10.1037/rev0000097>
- Zhu, M., Yang, Y., & Hsee, C. K. (2018). The mere urgency effect. *Journal of Consumer Research*, 45(3), 673–690. <https://doi.org/10.1093/jcr/ucy008>

Received November 19, 2020

Revision received January 23, 2023

Accepted February 22, 2023 ■