Journal of Behavioral Economics for Policy, Vol. 7, Issue 1, 25-33, 2023

Testing Behaviors in the Play of an Expected Prisoner's Dilemma

Richard Arend[1]*

Abstract
We test the impact of a new form of game structure – the expected game – focusing on whether it improves the outcomes of a prisoner's dilemma. An expected game is defined by the probability-weighted payoffs of two or more different reference or 'bookend' games. We analyze the impact of this structure and several relevant variants through a set of human subject-based experiments on non-naïve, consenting participants. As conjectured, we find that real play under this structure does decrease defection in a one-shot prisoner's dilemma game. Additionally, we find that such effects can be enhanced through framing and priming treatments. We discuss the micro-economic policy implications of our new results.

JEL Classification: C72; C7; C63; D81; L24

Keywords
implied games — expected payoffs — prisoner's dilemma — microeconomic policy

[1] *Professor of Strategic Management, University of Southern Maine*
***Corresponding author**: richard.arend@maine.edu

Introduction

Many policy-relevant economic and social interactions unfortunately entail the inefficiencies characterized by a prisoner's dilemma game (Axelrod, 1984). Thus, a significant amount of research has been devoted to analyzing modifications and structures that improve the outcomes of those interactions by increasing players' incentives to cooperate, through matching and re-matching protocols in repeated interactions (Andreoni & Miller, 1993; Yang et al., 2007; Duffy & Ocks, 2009), verbal communication (Tullock, 1999), and by letting players bargain over the game strategies in super-game (see the prisoner's dilemma game with confirmed proposals in Attanasi et al., 2013; 2015). One such new structure that has been argued to improve the outcomes is termed the *expected game* – where parties make their decisions based on the expected payoffs implied by two or more reference (*aka* bookend) games that each represent a different possible future for the players involved (Arend, 2020). We test the efficacy of the expected prisoner's dilemma game here, in addition to its manipulability through framing. We do so through a survey-based experiment with real, non-naïve decision-makers. The results provide a basis for improved policies that should reduce the inefficiencies of prisoner's dilemma-like (henceforth PD) situations.[1]

We accessed our experiment's participants through Amazon Turk and then filtered them for relevant game-playing skills. They played multiple games, including a standard PD and two versions of the expected PD, in one of four different treatment combinations. Upon analyzing the data, we found that the new expected game structure did improve outcomes (e.g., decreasing the defection rates of one-shot PD games). We also found that gain-loss framing and contextual priming treatments further enhanced those outcomes in the predicted directions. We discuss the relevant limitations and implications for policy after we explain how we arrived at our results.

Background

Arend (2020) recently introduced a structure for the expected game. This structure consists of at least two bookend reference games that define the decision the players face. The expected game's normal form payoff matrix is implied by the payoffs of the bookend games – each cell in the expected game's matrix is made up of the probability-weighted sum of the payoffs of the bookend games, where the weights add

[1] The PD model originated in work at the RAND Corporation in 1950 and was formalized shortly thereafter as a game involving prisoner plea-bargaining (Poundstone, 1993). Like many other simple games (e.g., chicken; matching pennies) that each represent a wide range of social interactions of interest (i.e., with the PD, the interest lies in gaining Pareto-improvements in the face of an individually rational choice not to), the PD has garnered attention and been examined in order to find better ways to play it and to modify it. Variants of the original one-shot PD include the finitely and infinitely repeated PD, the repeated PD with exit option, and where reputation can be tracked, where discount rates can apply, and where payoff manipulation by third-parties can occur. The repeated PD has been of great interest because the way it is best played theoretically can differ from how human subjects play it and the way a winning algorithm in a simulation plays it (Axelrod, 1984). Economics has used the PD to model many policy-related issues from strategic alliances to international trade to global warming mitigation. Our contribution is to test the effects of a new form of the PD – the expected PD – on human decision-makers, and to identify further treatments using that new form, in order to see whether defection can be reduced and cooperation increased. Such work adds to the literature, including the provision of some new (and newly-tested) tools, on how to address PD problems.

up to the full probability space.[2] To be clear, the bookend games are simply the set of games entailing the same choices and players as the expected game that they collectively generate through the weighted summing of the payoffs in each game cell. The bookend games depict alternative futures (e.g., where technical standard A wins versus where standard B wins in a current standards war, or where one future is defined by a piece of legislation passing and the other future where it does not) that involve different payoffs to the players based on their choices. See Figures 1–6 for examples of the standard and expected PD games: Figures 1–3 involve possible negative payoffs while Figures 4–6 involve only non-negative possible payoffs for the PD and for the expected PD games as defined by their sets of two bookends each.

RIVAL / YOU	L	R
U	1, 1	-2, 3
D	3, -2	-1, -1

RIVAL / YOU	L	R
U	-2, -2	-2, 8
D	8, -2	4, 4

RIVAL / YOU	L	R
U	4, 4	-2, -2
D	-2, -2	-6, -6

RIVAL / YOU	L	R
U	2, 2	-2, 1
D	1, -2	0, 0

RIVAL / YOU	L	R
U	0, 0	-2, 5
D	5, -2	-2, -2

Figures 1, 2 and 3 - the Standard PD, and the two Expected PD Games' Bookend Games (Depicted for the Treatment Involving Negative Possible Payoffs)

INSTRUCTIONS for Expected PD Games: You are playing a GAME composed of the TWO GAMES depicted. A fair coin toss will identify which of the two games provides the payoffs. You and your rival will NOT know the outcome of the coin toss prior to having to make your choice. So, given each game has an equal (50%) chance of determining the payoff outcome, what is your choice? Again, assume that you and your rival choose a best option without knowing what the other has chosen. [Choices are labeled U, D, and M – M being a mix of U and D, such as playing U based on a coin-flip.]

RIVAL / YOU	L	R
U	3, 3	0, 5
D	5, 0	1, 1

RIVAL / YOU	L	R
U	0, 0	0, 7
D	7, 0	2, 2

RIVAL / YOU	L	R
U	6, 6	0, 3
D	3, 0	0, 0

RIVAL / YOU	L	R
U	4, 4	0, 3
D	3, 0	2, 2

RIVAL / YOU	L	R
U	2, 2	0, 7
D	7, 0	0, 0

Figures 4, 5 and 6 - the Standard PD, and the two Expected PD Games' Bookend Games (Depicted for the Treatment Involving Non-Negative Possible Payoffs)

INSTRUCTIONS for Expected PD Games: You are playing a GAME composed of the TWO GAMES depicted. A fair coin toss will identify which of the two games provides the payoffs. You and your rival will NOT know the outcome of the coin toss prior to having to make your choice. So, given each game has an equal (50%) chance of determining the payoff outcome, what is your choice? Again, assume that you and your rival choose a best option without knowing what the other has chosen. [Choices are labeled U, D, and M – M being a mix of U and D, such as playing U based on a coin-flip.]

Game-theoretic reasoning supports improved outcomes under the expected PD structure – specifically, that when at least one of the bookends does *not* entail a defection-dominant strategy, then players can better rationalize choosing something other than defection. For example, when cooperation is the *maximin* choice across the bookend games, then players can rationally choose cooperation in the expected PD. Or, when cooperation is the dominant strategy in one at least one of the bookends, then players may be better able to justify at least a mixed approach to the expected PD. Arend (2020) mentions several bookend game characteristics that are more likely to induce cooperation, some of which we consider as treatments in our experimental test.

Our intention here is to test whether the game-theoretic reasoning underlying the predictions of outcome improvements in the expected PD structure translate into the real world – i.e., where boundedly-rational behaviors are at play. Prior behavioral economics studies have proven that real decision-makers, choosing among alternatives in an experimental setting, will *not* play with coldly rational optimization strategies (e.g., Axelrod, 1984; Sabater-Grande & Georgantzís, 2002; Simon et al., 1992). However, they will play, on average, in boundedly-rational ways that produce good payoffs; and, in expanding-pie games, often with better payoffs than when

[2] An *implied* game is determined by calculating the probability-weighted payoffs as p * left-side-game-payoffs + (1 – p) * right-side-game-payoffs for each cell of matrix (e.g., in Figures 2 and 3), where here the value of $p = ½$ is common knowledge.

playing fully rational strategies. As such, our expectations in running our experiment are consistent with the idea that the majority of participants will display rationality (rather than randomness) in their choices. And, that extends to a level of rationality that will support the predictions relating to the expected PD game. We believe that the results of human subjects experiments like the one we report on here represent legitimate behaviors and internal logics (as the myriad of similar decision-theory and game-theory testing research conducted over the past decades has proven – e.g., Binmore, 1994; Gintis, 2014; Kahneman & Tversky, 1979; Schram, 2005; Tversky & Kahneman, 1974; and Wakker 2010).

Given that our primary belief is that the expected structure will improve outcomes of the PD game, a natural next step involves testing whether those outcomes can be further enhanced with selected behavioral treatments. As such, we consider whether framing (e.g., Bargh & Chartrand, 2000; De Heus et al., 2010; Kahneman & Tversky, 2013; McDaniel & Sistrunk, 1991; Tversky & Kahneman, 1974) – can further decease defection and increase cooperation in the expected PD games. Specifically, we test whether gain-loss framing (e.g., De Dreu & McCusker, 1997), priming (e.g., Ahmed & Salas, 2011), and complexity (e.g., Hristova & Grinberg, 2010) affect behaviors in the playing of the game.

Data and Methods

We used *Amazon's Mechanical Turk* service (henceforth *MTurk*) to access, filter and pay our experimental game participants. We restricted the sample to those *MTurk* subjects with more experience (i.e., with 5000 or more tasks completed), who were rated as higher quality (i.e., with 97% and greater approval ratings), and who were more likely to understand multi-step instructions in English (i.e., those accessing *MTurk* from the US, UK and other high-English-fluency EU and Island nations). The full sample of 200 participants was obtained over three waves (i.e., of sizes 10, 100 and 90) posted within a ten-day period from late December, 2020 to early January, 2021. Two filter questions were used in order to only allow those respondents displaying a basic understanding of game theory (i.e., who were able to identify a dominant strategy in a three-choice, two-player normal-form game and able to calculate an expected value of a simple bet) to partake in the main survey, where they had to answer all questions in order to be paid.[3] Monetary incentives were designed as follows: Each subject received a posted rate of compensation for completing the survey; in addition, participants were notified that one of the survey questions – in the form of a game – would be used as a basis for extra compensation (i.e., a bonus for solving that game). Average payment per subject was $8.92, for an average duration of the online experiment of approximately 22 minutes.

The survey included a preliminary set of demographic questions encompassing several decision-maker characteristics. Relevant characteristics included: *male* (yes/no) – 65%; *age* – 38 years old on average; *US-id* (identify as American – yes/no) – 92%; *college* (education level at college undergraduate or better – yes/no) – 77%; *exposure to game theory* (5-point Likert scale from 1 = none to 5 = expert) – 2.2 average; and, *experience calculating expected values* (5-point Likert scale from 1 = none to 5 = a great deal) – 2.6 average. Thus, our survey participant pool skewed male, mature, American, educated and of moderate relevant knowledge of logics and mathematics. These characteristics were used in robustness checks (i.e., by controlling for such characteristics when regressing the influence of a treatment on a focal outcome, like the choice of defection) because such characteristics have been known to influence strategic choices (e.g., Francioni et al., 2015).

After our experiment's participants successfully answered the two filter questions, they consented to the main study, and provided the demographic information described above. The main study involved each participant playing a sequence of 2x2 normal form games (i.e., two-player games where each player has two choices, where choices are made simultaneously, and where the payoffs are symmetric). That sequence of games included: the standard PD; two bookend games (one cooperation-dominant, and the other defection-dominant), immediately followed by the expected PD composed of an equally-weighted probability of those bookends; and, an expected PD composed of an equally-weighted probability of different bookends (one bookend with a mixed-strategy Nash equilibrium, and the other with a defection-dominant strategy). All three PD games – the standard PD and the two expected PDs – entailed the same expected payoff matrix for the participant. In each game, participants could choose their strategy from three options – two being pure choices (i.e., labeled *U* [for up] and *D* [for down]) and one being an unspecified mix of the two pure choices (i.e., labeled *M* [for mixed]). All games were played against a hypothetical rival.[4] Immediately after playing each PD game in the sequence, participants were asked to identify their primary rationale for their choice from the following set: it was the dominant strategy; it maximized the expected value; it maximized the minimum payoff; it maximized the maximum payoff; it is my best guess and feels right; and, I don't know.

[3] As with most surveys that draw from the *MTurk* respondent pool, we paid participants only when they answered all of the questions and indicated the expected effort in doing so (e.g., by checking that they did not rush through the survey, they did not answer all questions with the last option, and so on); as *MTurk* regular participants, they know that they have an incentive to put effort into their answers because if they are found out and not paid then their reputation suffers, decreasing the likelihood that they qualify for future well-paying assignments.

[4] Subjects were notified that they were playing against a hypothetical rival in every one-shot game. To this subject pool – *MTurk* respondents making choices in a survey asynchronously – that implied they were simply trying to choose, and then justify, their best strategies. Because they had passed the filter questions, they knew that they were not playing against another person, nor against a computer. The outcome of any game was never provided; to be clear, subjects were always instructed to identify the best option (of the three possible) in each simultaneous-move game.

Test and Treatments

While the main test involved determining whether the expected PD game improves outcomes over the standard PD game, we also tested *four treatment manipulations* in order to determine whether that main effect could be enhanced in specific ways. The first treatment manipulation involved *gain-loss framing* (i.e., where payoffs either did or did not display possible losses) – here, the PD games played either entailed only non-negative possible payoffs in all games (PDs and bookends) or included negative possible payoffs. The second treatment manipulation involved *priming* – here, the sequence either entailed playing the standard PD prior to or after playing the first expected PD. Roughly one-quarter of the participants played each of the four combinations of these two treatment variables (gain-loss framing and priming). All participants experienced a third treatment manipulation. That treatment manipulation involved playing a *more complex* expected PD game – where the first expected PD game was defined by two bookends, each with a different dominant strategy, the second more complex one involved a bookend game with a mixed strategy as the rational choice. The idea was to inquire whether a less-clear dichotomy of bookend games would affect play. A fourth treatment manipulation was experienced by roughly half of the participants. Those playing the games involving possible negative payoffs confronted a first set of bookends that had *cooperation as the clear maximin choice*. The idea was to see whether that rationalized choice was consciously selected more in those games versus in the other expected PD variants (denoted as *ePD*s henceforth).

Method of Testing

We report the results of proportions difference tests (two-tailed) for the main PD versus ePD test and for each of these treatment manipulations. We report the proportions of the treatment and the non-treatment samples, and whether those differed in a statistically significant way.[5]

Results

The main result is that the ePD structure works as argued in Arend (2020) – with outcomes improving over the standard PD. Specifically, defection drops significantly (from a proportion of 77% to one of 65%; $p < 0.01$). And while the direct choice of cooperation increases in a non-statistically significant way (from a proportion of 18% to one of 20%), the implied cooperation level (calculated by combining the direct choice of cooperation with its indirect choice – by assuming that half of the mixed strategy choices lead to cooperation) increases marginally significantly (from a proportion of 21% to one of 28%; $p < 0.06$). This is partly because the choice of mixed strategy increases significantly (from a proportion of 6% to one of 15%; $p < 0.001$). Table 1 summarizes the results of our study (including the main tests and the treatment-comparison tests, as well as providing information about the sampling). Figures 7-10 illustrate the results of the main test's different dimensions.

In order to better understand what could be driving the difference in outcomes between the PD and the ePDs, we considered the level of *rationally-consistent* choices made by the participants in each game type. To be considered a rationally-consistent choice required that the choice made among the three possible options was consistent with one of the four specific rationales (where the last two of the six possible rationales were excluded because each was non-specific). The level of rational consistency dropped significantly (from a proportion of 73% to one of 45%; $p < 0.001$). This drop was mostly driven by fewer choices of the dominant strategy and rationale when the ePDs were played. There are several possible explanations for this result. The explanation that is consistent with the arguments made about the ePD is that its structure works as intended, but with the consequence that players have a harder time identifying the proper justification for their rationally-chosen actions to defect less. An alternative explanation is that the impact of the new structure is caused by the added confusion it creates to the participant – as apparently indicated by the loss in their ability to justify their choices when confronted by the more complex structure (i.e., where they had to calculate the expected value payoffs of the implied game that they are playing). However, we do not believe that the confusion explanation is appropriate because several of the treatment results appear to indicate that participants are capable of dealing with the more calculation-intensive games. We comment on that assertion in the results that follow.

Treatment Results

For the first treatment manipulation – testing the impact of gain-loss framing – the only statistically significant effects arose in the second ePD game (i.e., the one with one mixed-strategy bookend). This treatment did not affect the outcomes in the first ePD game or in the standard PD game in a significant way. For the second ePD game, there was a significant drop in defection (from a proportion of 75% to one of 61%; $p < 0.05$) and a significant increase in pure cooperation (from a proportion of 15% to one of 28%; $p < 0.05$), when the game only involved non-negative possible payoffs. (There was no statistically significant change in the choice of a mixed strategy or in the rate of rationally consistent choices made.) The net effect of this testing appears to be a weak indication that the positive impact of the ePD game can be enhanced by framing the game as one where losses are not possible (e.g., perhaps by depicting the games only in terms of gross benefits).

For the second treatment manipulation – testing the impact of priming – the relevant effects relate only to the first ePD

[5] Note that we do not report on the testing of the influence of demographic characteristics because these were neither focal to the research question of this paper nor did they influence the main results. The regression analyses that controlled for the influence of demographic characteristics did not affect the treatments' effects reported, regardless if some of those factors entailed inconsistently significant but small absolute effects, in the expected directions (e.g., with males being more likely to choose defection).

Test/ Treatment	Numbers	Relevant [statistically-significant] Results
F-Testing the impact of the new form of game structure – the expected game – on the outcomes of the prisoner's dilemma (PD): does defection increase and cooperation increase in the expected PD (ePD)? [same subjects playing different games]	All playing PD (N_1=200); all 200 playing two ePDs (N_2=400).	• Pure-D↓ 77% to 65% proportions: z = 2.91, p = 0.004 • U-with-half-Mix↑ 20.5% to 27.7% proportions: z = 1.90, p = 0.06 • Pure-M↑ 6.0% to 15.3% proportions: z = 3.28, p = 0.001 • Rationale-Consistent-with-Choice↓ 73% to 45% proportions: z = 6.31, $p < 0.001$
Treatment Variable #1 = gain-loss framing. Do these games have different results with non-negative payoffs relative to when some possible payoffs are negative? [different subjects play gain versus loss games]	About half play gain frame (N_1=102); rest play loss frame (N_2=98).	*Only* for the second ePD – in terms of gain vs. loss frame: • Pure-D↓ 75% to 61% proportions: z = 2.07, p = 0.04 • Pure-U↑ 15% to 28% proportions: z = 2.25, p = 0.03
Treatment Variable #2 = priming. Does playing the PD before the ePD induce more defection (by priming that behavior)? [different subjects play before versus after sequence]	About half play ePD after PD (N_1=102); rest play ePD before PD (N_2=98).	In terms of PD before (priming) vs after: • Pure-D↑ 69.4% from 54.9% proportions: z = 2.11, p = 0.04 • Pure-U↓ 12.2% from 23.5% proportions: z = 2.08, p = 0.04
Treatment Variable #3 = ePD complexity. Does the complexity of the ePD (e.g., involving a mixed rather than dominant strategy bookend) affect the outcome (when the same payoffs of the combined bookends occur)? [same subjects playing different games]	All playing easier ePD (N_1=200) before harder ePD (N_2=200).	Comparing the harder to the easier ePD: • M↓ 10.5% from 20.0% proportions: z = 2.64, p = 0.01
Treatment Variable #4 = ePD maximin choice clarity. Does clarity of the maximin choice in an ePD (i.e., only one choice has maximizes the minimum possible payoff – with no ties in any conditions) generate greater rational choice of that strategy (i.e., where the choice is backed up by the correct rationale indicated)? [different subjects play gain versus loss games]	About half play unclear maximin choice (N_1=102); rest play clear choice (N_2=98).	Comparing the clear to the unclear ePD maximum condition: • Correctly-Rationalized- Choice↑ 7.1% from 2.0% proportions: z = 1.74, p = 0.09

List of *Non-Significant* Results [proportions (%) reported]: Test: pure-U changed from 17.5% to 20.0%; Treatment#1: for PD (pure-D unchanged at 76.5%; pure-U changed from 16% to 19%) and for first ePD (pure-D changed from 59% to 65%; pure-U changed from 17% to 19%); Treatment#3: pure-D changed from 62% to 68%; pure-U change from 17% to 15%. All proportions tests are two-tailed.

Table 1. Summary of Tests, Treatments and Empirical Results

game. (This is because the second ePD game was always played last in the sequence, and so was always primed by both a standard PD and an ePD.) When *not* priming with a standard PD game (compared to priming with it), defection rates decreased significantly (from a proportion of 69% to one of 55%; $p < 0.05$), and pure cooperation increased (from a proportion of 12% to one of 24%; $p < 0.05$). The choice of mixed strategy did not differ significantly. The net effect of testing appears to provide a strong indication that the positive impact of the ePD game can be enhanced by controlling the

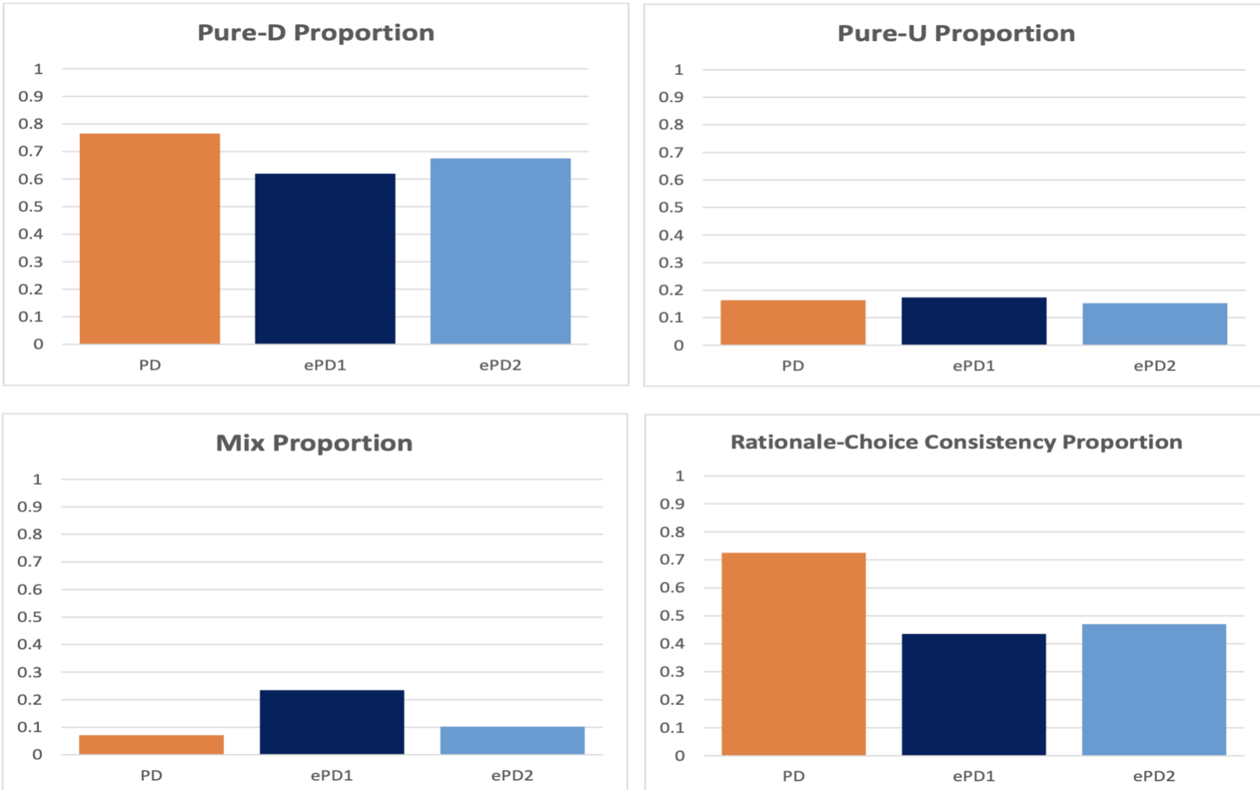

Figures 7-10 – Proportions Comparisons of Defection, Cooperation and Mixed Strategy Choices, and of the Rationale-Compatible-Choices, for the Standard and Expected Prisoner's Dilemma Games

priming – here, ensuring that a standard PD game is *not* played prior to the focal ePD game.

Interestingly, the rate of rationally-consistent choices also dropped significantly (from a proportion of 48% to one of 30%; $p < 0.05$) under this treatment. That is consistent with the result for the main test result where rationally-consistent choices also dropped significantly. But it should be noted that a similar drop in the second ePD game also occurred for this treatment (a drop from a proportion of 48% to one of 35%; $p < 0.10$). Priming appears to have helped rational consistency in both ePDs. The most obvious possible explanation involves momentum – when players are primed to defect based on a rational strategy (from engaging in the standard PD game first where *D* was the dominant choice) they are more likely to defect with an identifiable rationale throughout (i.e., in the two subsequent ePDs). Players without priming made different choices (significantly so in the first ePD) but struggled more with identifying their rationales. (This points to possible issues over identifying the rationale rather than any confusion over making a desired choice.)

The third treatment manipulation – involving the question of whether an ePD game with a non-dominant strategy bookend affected the outcomes – entailed one statistically noteworthy result. Comparing the outcomes of the first ePD game with those of the second revealed significant effects only for the prevalence of the mixed strategy choice. It dropped significantly (from a proportion of 20% to one of 11%; $p < 0.01$) while the proportions for both cooperation and defection increased, but insignificantly. This treatment does *not* appear to enhance the main ePD-related effect. Nevertheless, the one significant result implies that the strategy choice of 'mix' is more likely when the ePD's bookends entail opposing dominant optimal strategies. (This result speaks to the rationality of the participants; here, the result implies that they recognize the possibility of each bookend determining the final payoffs, and then act so as not to miss out on the opportunity to have played the cooperation-dominant bookend when it is equally likely to occur. Again, this appears to go against the idea that participants were confused by the ePD.)

The fourth treatment manipulation – whether an ePD game with a clear maximin choice of cooperation would increase the occurrence of that rationally-consistent choice – had notable results. The treatment revealed a marginally significant increase of the choice of cooperation-with-the-correct-identification-of-the-maximin-rationale (from a proportion of 2% to one of over 7%; $p < 0.10$). While still an uncommon choice-rationale combination (occurring less than 10% of the time in that treatment but explaining over 40% of the pure cooperation choices), the result was consistent with the expected effect of that treatment.

Discussion and Conclusions

We set out to test the effects of a new structure on the PD game – a structure that has been argued to decrease defection and increase cooperation. In addition to testing that structure's effects directly, we also tested selected treatment manipulations on that new structure to determine whether such positive effects could be enhanced. Our human-subjects experiment provided several statistically-significant results that supported those main predicted impacts in addition to the use of several framing treatments that appeared to enhance those impacts.

Policy Implications

Solving a PD problem – in terms of improving cooperation – is an established stream in the social science literature and one that has historically been framed as a policy challenge to institutions (e.g., to governments, to organized religions, and to other forms of social contract – Axelrod, 1984; Freebairn, 2008; Shughart & Tollison, 2005). Because so many human interactions can be modeled with the PD payoff form, it is a game that is analyzed in many social sciences including business, economics, politics and sociology. In business, it has been used to model, among other interactions, strategic alliance behavior (e.g., Arend & Seale, 2005). In economics, it has served as a model for a range of interactions from oligopolistic competition to collective action that produces a collective good (e.g., Axelrod, 1980; Georgantzís & Attanasi 2016). In politics, it has served as a model for arms races and other conflicts (e.g., Majeski, 1984). And, in sociology, it has served as a model for multi-player interactions like the tragedy of the commons (e.g., Nicholson, 2000), as well as for specific issues in societal contests like with doping in sport (e.g., Schneier, 2012). We see policy-makers having a positive role in adjusting any such 'given PD' in ways to improve social welfare, where one of those adjustment options may be in transforming the given PD into an ePD in order to increase cooperative actions. The idea that policy-makers can step in to adjust given games has been standard in economics for decades (e.g., in international trade – see Brander & Spencer, 1985) and so should sometimes apply in these cases as well.

There are four primary ways that knowledge of the ePD – from our study's results – could benefit the policy-maker: First, when the policy-maker is a primary party in an ePD that is generated because alternative futures exist in possible demand or technology conditions, or in regulations, or in natural or man-made disasters, for example, our results provide insights into how to interpret the inferred (i.e., expected) game and play it 'better' (i.e., either by exploiting the other party or by mutually finding greater cooperation and higher Pareto efficiency). Second, when the policy-maker is a third party in a PD, then our results provide an option for altering it into an ePD that could produce greater social benefits. For example, the policy-maker could attempt to change the given context from a one-clear-future perception into a two-alternative-futures perception by introducing legislation to provide that uncertain alternative (e.g., as a possible new tax, subsidy, tariff or regulation that passes and is enforced in the alternative future). Third, when the policy-maker is an active or passive creator of a PD (e.g., Clark & Lee, 2003), it may wish to help design it to minimize damage by using our results to convert it into an ePD instead (e.g., by introducing possible new policies that open up the future to more than one clear state). And, fourth, when even the ePD is not expected to provide sufficient cooperation (based on our results), the policy-maker (as the government) may simply wishes to directly intervene and enforce players' choices (i.e., to collectivize decision-making through the state – Orbell & Wilson, 1978).

The results from our treatment testing provide several further implications for policy aimed at improving the outcomes of multi-party interactions that could be represented in a PD-like form. For example, if policy can be used to avoid priming the parties (i.e., to avoid an immediately prior standard PD experience), or if policy can frame the payoffs to avoid negative possible payoffs (e.g., by quoting gross rather than net payoffs), or if policy can structure the bookend games so that cooperation is the maximin choice, then the beneficial impact of introducing an ePD structure is likely to be further improved. These preliminary results serve as a reminder to policy-makers to drill down on PD situations in order to reveal the extent of the underlying relationships of inputs to outcomes, especially when those are conditional or probabilistic (i.e., when they could involve an ePD), because that may provide the opportunity to exploit those conditions in new structures as we have shown here.

Limitations

We admit several limitations on the generalization of our study's results. Our experiment was limited by: the number of participants;[6] the use of *MTurk* subjects (rather than a more random sample); the demographics of those subjects who made it through our filtering (i.e., subjects who skewed more male, mature, US-based and college-educated); the relatively low stakes involved; and, the number of treatments applied. That said, our main results appeared robust (i.e., when regressing on the outcomes while controlling for demographic factors).

Future Work

Besides further studies that could address some of the limitations described above, there are several interesting areas for future work to explore. Our results suggest future work

[6]As with most experimental empirical work, we were limited in sample sizes used for testing (e.g., due to budgets and logistics). This potentially affected some of the power of our results. Power analyses indicate that (for 80% power and 5% type-1 error) the sample sizes for the main test and for the first two treatment manipulations were appropriate, but that the size for the third treatment manipulation was marginally low and the size for the fourth treatment manipulation was significantly lower than recommended (due to the low proportion level at play).

should consider testing further treatment manipulations to identify more of them that enhance cooperation in the ePD; it should also aim to more fully tease out why participants alter behaviors under new environments (e.g., whether it is based on rationality or a heuristic or something else). Regardless, we do need to continue to test such newly proposed solutions to policy problems that entail PD-like payoff structures not only because they can lead to improvements in social welfare but also because such testing itself often produces further new areas to explore to enhance those impacts.

Declarations and Ethics Statements

Ethical approval – IRB-approved.
Informed consent from participants – yes (as part of the IRM approval).
Data availability statement for Basic Data Sharing Policy – upon request (by journal for review or by readers if published).
Competing interests – none.
Funding – none specific.

References

Ahmed, A.M. & Salas, O. (2011). Implicit influences of Christian religious representations on dictator and prisoner's dilemma game decisions. *The Journal of Socio-Economics, 40*(3), 242-246.

Andreoni, J., & Miller, J. H. (1993). Rational cooperation in the finitely repeated Prisoner's Dilemma: Experimental evidence. *Economic Journal, 103*, 570–585.

Arend, R.J. & Seale, D.A. (2005). Modeling alliance activity: an iterated prisoners' dilemma with exit option. *Strategic Management Journal, 26*(11), 1057-1074.

Arend, R.J. (2020). The expected prisoner's dilemma with rationally arising cooperation. *PLOS ONE, 15*(9), e0239 299.

Attanasi, G., García-Gallego, A., Georgantzís, N., & Montesano, A. (2013). An experiment on prisoner's dilemma with confirmed proposals. *Organizational Behavior and Human Decision Processes, 120*(2), 216-227.

Attanasi, G., García-Gallego, A., Georgantzís, N., & Montesano, A. (2015). Bargaining over strategies of non-cooperative games. *Games, 6*(3), 273-298.

Axelrod, R. (1980). Effective Choice in the Prisoner's Dilemma. *The Journal of Conflict Resolution, 24*(1), 3–25.

Axelrod, R. (1984). The Evolution of Cooperation. Basic Books; New York.

Bargh, J.A. & Chartrand, T.L. (2000). "The mind in the middle: A practical guide to priming and automaticity research". In *Handbook of research methods in social and personality psychology*; Judd, H.T. & Reis, M.C., Eds.; Cambridge University Press: New York, NY; pp. 253–285.

Binmore, K.G. (1994). *Game theory and the social contract: just playing.* (Vol. 2). Cambridge, MA: MIT Press.

Brander, J.A. & Spencer, B.J. (1985). Export subsidies and international market share rivalry. *Journal of International Economics, 18*(1-2), 83-100.

Clark, J.R. & Lee, D.R. (2003). The Increasing Difficulty of Reversing Government Growth: A Prisoners' Dilemma that Gets Worse with Time. *Journal of Public Finance and Public Choice, 21*(2-3), 153-165.

De Dreu, C.K.W. & McCusker, C. (1997). Gain–Loss Frames and Cooperation in Two-Person Social Dilemmas: A Transformational Analysis. *Journal of Personality and Social Psychology, 72*(5): 1093–1106.

De Heus, P., Hoogervorst, N. & Van Dijk, E. (2010). Framing prisoners and chickens: Valence effects in the prisoner's dilemma and the chicken game. *Journal of Experimental Social Psychology, 46*(5), 736-742.

Duffy, J., & Ocks, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior, 66*, 785–812.

Francioni, B., Musso, F. & Cioppi, M. (2015). Decision-maker characteristics and international decisions for SMEs. *Management Decision, 53*(10), 2226-2249.

Freebairn, J. (2008). Some sectoral and global distributional issues in greenhouse gas policy design. *Agenda: A Journal of Policy Analysis and Reform, 15*(1), 13-28.

Georgantzís, N., & Attanasi, G. (2016). *Non-competitive markets. In Experimental Economics* (pp. 21-36). Palgrave Macmillan, London.

Gintis, H. (2014). *The bounds of reason.* Princeton, NJ: Princeton University Press.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263-291.

Kahneman, D. & Tversky, A. (2013). "Prospect theory: An analysis of decision under risk". In *Handbook of the fundamentals of financial decision-making: Part I* (pp. 99-127).

Majeski, S.J. (1984). Arms races as iterated prisoner's dilemma games. *Mathematical Social Sciences, 7*(3), 253–266.

McDaniel, W.C. & Sistrunk, F. (1991). Management dilemmas and decisions: Impact of framing and anticipated responses. *Journal of Conflict Resolution, 35*(1), 21-42.

Nicholson, W. (2000). *Intermediate microeconomics and its application* (8th ed.). Fort Worth, TX: Dryden Press: Harcourt College Publishers.

Orbell, J.M. & Wilson, L.A. (1978). Institutional solutions to the N-prisoners' dilemma. *American Political Science Review, 72*(2), 411-421.

Poundstone, W. (1993). *Prisoner's Dilemma.* New York: Anchor.

Sabater-Grande, G., & Georgantzís, N. (2002). Accounting for risk aversion in repeated prisoners' dilemma games: An experimental test. *Journal of Economic Behavior & Organization, 48*(1), 37-50.

Schneier, B. (2012). Lance Armstrong and the Prisoners' Dilemma of Doping in Professional Sports — Wired Opinion. *Wired.* Wired.com (retrieved 10-29-2012).

Schram, A. (2005). Artificiality: The Tension between Internal and External Validity in Economic Experiments. *Journal of Economic Methodology, 12*(2): 225–37.

Shughart, W.F. & Tollison, R.D. (2005). Public choice in the new century. *Public Choice, 124*(1), 1-18.

Simon, H., Egidi, M., Marris, R. L., & Viale, R. (Eds.). (1992). *Economics, bounded rationality and the cognitive revolution.* Edward Elgar Publishing.

Tullock, G. (1999). Non-Prisoner's Dilemma. *Journal of Economic Behavior and Organization, 39*, 455–458.

Tversky, A. & Kahneman, D. (1974). *Judgment under uncertainty: Heuristics and biases. Science, 185*(4157), 1124-1131.

Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity.* Cambridge University Press.

Yang, C.-L., Yue, C.-S. J., & Yu, I.-T. (2007). The rise of cooperation in correlated matching Prisoner's Dilemma: An experiment. *Experimental Economics, 10*, 3–20.