Contingent Reasoning and Dynamic Public Goods Provision*

Evan M. Calford[†] and Timothy N. Cason[‡]

April 9, 2021

Abstract

Individuals often possess private information about the common value of a public good. Their contributions toward funding the public good can therefore reveal information that is useful to others who are considering their own contributions. This experiment compares static (simultaneous) and dynamic (sequential) contribution decisions to determine how hypothetical contingent reasoning differs in dynamic decisions. The timing of individuals' sequential contribution decisions is endogenous. Funding the public good is more efficient with dynamic than static decisions in equilibrium, but this requires the decision-maker to understand that in the future they can learn from past events. Our results indicate that a large fraction of subjects appreciate the benefits of deferring choice to learn about the contribution decisions of others. The bias away from rational choices is in the direction of Cursed equilibrium on average, particularly in the static treatment.

Keywords: Cursed equilibrium; Voluntary contributions; Club goods; Laboratory experiment

JEL Codes: C91, D71, D91, H41

* *Preliminary draft. Please do not circulate without permission.* We thank William Brown for excellent research assistance, and audiences at the Economic Science Association conference for valuable comments.

† John Mitchell Fellow, Research School of Economics, Australian National University, Canberra, Australia. Evan.Calford@anu.edu.au

‡ Krannert School of Management, Purdue University, Indiana, USA. cason@purdue.edu

1 Introduction

Collective action requires coordination and often involves uncertainty. Besides strategic uncertainty about other agents' behavior, in many realistic situations the value of taking collective action is unknown until it takes place. Examples range from global challenges such as climate change, to specific goals addressed in thousands of crowdfunding campaigns that support civic objectives, create art or develop new products. In many, or perhaps even most of these problems, the uncertain value has a strong common value component that correlates individuals' benefits from public good provision. Moreover, contributions to support collective action are typically not simultaneous, as firms, individuals and governments maneuver over time to address funding needs or implement regulations to accomplish common goals.

This dynamic process of common value public goods provision can lead to progressive revelation of agents' private information about the value of collective action. Their ability to make use of this information to provide only the most valuable public goods has important welfare consequences, but in naturally-occurring economic environments the beliefs, signals and preferences of individuals are typically unobservable. This creates difficult challenges for empirical research on collective action. In this paper we present a model comparing dynamic and static contributions to common value public goods. We also report a laboratory experiment examining predictions of the model arising for fully rational agents who correctly condition on private and public information, as well boundedly rational agents who have difficulty with contingent reasoning. Our data suggest a mixture of rational and boundedly rational types, with subjects having an easier time with contingent reasoning in the dynamic treatment.

The implications of limited statistical reasoning by humans, and particularly their failure to understand how others' actions provide valuable signals of their own private information, has been documented beginning with early evidence of the "winner's curse" in common value auctions (Capen et al., 1971; Kagel and Levin, 1986, 2002). This was formalized by Eyster and Rabin (2005) for more general environments, who introduced the notion of "cursed equilibrium." In a cursed equilibrium for common value auctions, bidders best respond to incorrect beliefs that fail to account for how rival bidders' bids depend on their signals (for a survey see Eyster (2019)). Robust evidence of this type of limited rationality has been provided in a range of environments, from simplified nonstrategic settings such as the "Acquire a Company" problem (Bazerman and Samuelson, 1983; Charness and Levin, 2009) to voting (Esponda and Vespa, 2014) and non-auction market environments (Ngangoue and Weizsacker, 2021). Few previous studies have explored how limitations for contingent reasoning affect choices in a common value public good setting (Cox, 2015).

Reasoning about hypothetical contingencies is difficult for humans, and the implications of such reasoning failures are profound. Esponda and Vespa (2014) found that subjects in their voting experiment were much better at drawing inferences about actual previous decisions of others, available only in a sequential setting, than hypothetical events needed to guide choices in a static environment. For the static setting with simultaneous decisions, a majority of subjects behaved nonstrategically, in contrast to the sequential setting where many were able to extract information from the observed voting decisions of others. This motivates our comparison of dynamic and static public goods provision in the present paper, with the overarching goal to provide more insight into the source of difficulties people have with contingent reasoning. As summarized in Table 1, and in contrast to previous studies in which the sequential ordering is enforced, in our dynamic environment the agents choose the timing of their own decisions. That is, in our public goods setting, an agent in our Dynamic treatment can elect to contribute to the public good in the first period or may elect to delay the decision until a later period.

	$T = 1$	$T = 2$
Static (simultaneous)	P1, P2	
Sequential	P1	P2
Dynamic	P1, P2	P1, P2

Table 1: Order of decision making under various timing structures, where $P1$ denotes player 1 and $P2$ denotes player 2. Previous studies on information extraction and contingent reasoning have focused on comparing player 2 behavior across equivalent Static and Sequential environments. Here, we compare behavior across Static and Dynamic environments.

The failure of contingent reasoning regarding hypothetical, but not observed, states of the world is not well understood. Here, we posit two possible sources for this failure. One possibility is that individuals simply fail to recognize that there is information in others' simultaneous decisions that they could find useful for their own judgments and belief updating; we refer to this type of naivete as *unawareness*. It reflects the systematic disregard that others' relevant private information could be partially revealed through their hypothetical actions. However, after being informed of the actions of others, the arrival of the information alerts the individual to the belief updating problem. Alternatively, subjects may be aware that information exists that could be extracted and useful, but they have difficulty doing so when reasoning must be hypothetical. We call this *complexity*. After being informed of the actions of others the information extraction problem is simplified, such that the individual may now respond effectively.

Several studies have shown how participants are better able to exhibit contingent reasoning when information is not hypothetical, for example by making choices sequential so that payoff consequences of each contingency are more transparent (Esponda and Vespa, 2014, 2019; Ngangoue and Weizsacker, 2021) or by reducing the underlying uncertainty in the environment (Martinez-Marguina et al., 2019). In these previous papers, unawareness and complexity are confounded: the introduction of information both simplifies the choice environment and brings attention to the information extraction problem.

The dynamic structure of our experimental design allows for a partial separation of the unawareness and complexity explanations for the failure of contingent reasoning. First, note that a comparison of the second stage ($T = 2$) of our dynamic treatment with the static treatment is analogous to the comparison in the previous literature: information arrives before the second stage begins, and this information both simplifies the choice problem and highlights the existence of the information. Importantly, we can also compare behavior in the first stage of the dynamic treatment with behavior in the static treatment. If *complexity* is the source of contingent reasoning failures, then we should expect that subjects prefer to delay decision making to future stages where the arrival of information will make the decision less complex. However, if *unawareness* is the source of contingent reasoning failures then the subjects will be completely ignorant about the value of waiting for information to arrive and should therefore behave identically across the static treatment and the first stage of the dynamic treatment.[1]

Recently, Li (2017) has studied complexity in the context of mechanism design, where a strategy is considered to be simple (i.e. not complex) if the lowest payoff it provides is greater than the largest payoff of any other strategy. Oprea (2020) has studied complexity as it relates to the implementation of strategies where, among other conditions, the implementation of a strategy is perceived as more complex the greater the number of states it must be conditioned on. By that definition the dynamic environment is more complex than the static setting because it has more states. Our notion of complexity is distinct from each of these, and is a complexity of calculation. We suggest that, even when a person knows that they should condition a calculation on particular state(s) of the world, the mere existence, and potential realization, of non-payoff relevant states make the calculations more difficult to perform.

[1] As is standard in experimental economics, our experimental subjects play the same game multiple times. Thus, we might expect that unawareness diminishes over time as subjects learn about the strategic environment. Our data exhibits only very weak learning effects, however, suggesting that either unawareness never existed among the population or that the unawareness was of a type that could persist through multiple repetitions of the game.

Our experiment also addresses new issues in information extraction and contingent reasoning. In previous studies comparing simultaneous and sequential choices to explore contingent reasoning, the outcomes are isomorphic in the sense that optimal choices and equilibrium outcomes do not vary when introducing the sequential game form. Esponda and Vespa (2019), for example, effectively change the framing of the decision tasks on five classic problems, helping subjects focus on the set of states where their choice matters. This is not the case for our public goods provision problem, which features a more complex signal space, greater payoff uncertainty, and a richer dynamic structure.[2] This confronts decision makers with a novel information extraction problem that previous research comparing sequential and simultaneous contingent reasoning experiments has not addressed. Here, agents have an option value from deferring their decision about whether to contribute to the public good. The option value of waiting leads some to initially delay support, which they must trade off against the opportunity to signal to others their favorable information about the good's common value. Effective contingent reasoning leads to information sharing across multiple decision rounds in the dynamic treatment that more efficiently reveals information through the "richer" message space. The public good is provided less frequently in equilibrium in the dynamic than the static treatment, with a pronounced drop in provision when it has a low common value and should not be provided. More information is revealed about the value public good through the multiple rounds of choices made in the dynamic treatment.

The option value of waiting consists of two components in our setting. First, the value of flexibility and, second, the value of information. On the other hand, there is also a value of commitment for an agent with a favorable signal about the value of the public good. By committing to the public good in the first stage, the agent can credibly signal her private information to others. As a consequence, in the dynamic treatment, there is more than one dimension in which an agent can have cursed beliefs. The agent might believe that others do not condition behavior on their private signals (the standard, or first order, cursedness problem), or the agent might believe that other players do not condition behavior on the observed behavior of others (which we call the second order cursedness problem). First order cursedness decreases the option value of waiting, while second order cursedness decreases the value of commitment.

Our results indicate that a large fraction of subjects appreciate the benefits of deferring choice to learn about the contribution decisions of others when their signals about the public good value are near the margin. They also react to the information conveyed by others' choices, and how

[2] Multiple equilibria exist in all of our treatments, which also raises interesting new questions about behavioral equilibrium selection with contingent reasoning. For our empirical analysis we focus on symmetric equilibria in which the public good is provided with positive probability.

others' choice to select the public good signals a higher common value. The bias away from Nash equilibrium choices is in the direction of Cursed equilibrium on average, particularly in the static treatment. Overall, however, public good provision rates and errors in overprovision do not differ in the static and dynamic treatments as predicted by equilibrium. That is, while there is substantially less commitment to the public good in the first stage of the dynamic treatment than the static treatment, the aggregate provision rate in the dynamic treatment is increased to static levels via additional commitment opportunities the later stages.

2 Experimental Design

Several experimental papers have considered uncertainty in social dilemmas, including uncertain returns to contribution.[3] Our design features two key stylistic departures from the standard paradigm of public goods contribution games. Under standard game theoretic assumptions of risk neutral expected utility agents, our design changes have no effect on expected behavior. However, for subjects who exhibit risk aversion or status quo bias, our design deliberately ameliorates the potential behavioral impacts of deviations from this theoretical benchmark. This allows us to better isolate the role of contingent thinking in public good provision, which is of principal interest.

Consider first our static treatment. Three subjects must simultaneously decide whether they prefer a common value public good (PG) or not. The PG is provisioned only if at least two of the three subjects state a preference for the PG. Further, the PG is excludable; if only two subjects state a preference for the PG, then the third subject is excluded from the benefits of the PG.[4] We contrast our design with the typical approach where expressing a preference for the PG is operationalized as a requirement to pay a monetary contribution to the public good. For example Cox (2015), which studies an environment that is very close to our static treatment, includes a refund mechanism such that payments are refunded if the funding threshold for the PG is not met. This mirrors the way that leading crowdfunding sites such as Kickstarter operate: pledges to purchase a good or support a public good only result in payments if sufficient pledges by others

[3]Some of these studies find that contributions are lower with uncertain public returns (Dickinson, 1998; Gangadharan and Nemes, 2009; Levati et al., 2009), while others do not indicate contribution impacts of uncertainty, such as Stoddard (2017). Few studies have considered uncertain returns to common-value public goods, other than Cox (2015). See Cox and Stoddard (2021) for further discussion, and a static public goods provision experiment with information sharing about public returns through (binary) cheap talk messages.

[4]It is, perhaps, more accurate to refer to the good here as a club good rather than a public good. Nevertheless, we follow convention in the recent experimental literature and use the term public good throughout.

bring the total to the required threshold.

In the standard design, not paying for the PG is therefore the default action and this could, potentially, promote a status quo bias against the provision of the PG. In contrast, our design operationalizes expressing a preference for the PG as a binary choice between the PG and a private good, so it does not have a status quo: each subject must make an active choice between the PG and private good. Relatedly, this reframing of the problem removes any concerns regarding the potential for a wedge between the willingness-to-pay and willingness-to-accept for the PG.

The PG has a common value, with each subject receiving a private signal that is partially informative of the true value. Therefore, in the standard design, where subjects pay into the PG, the PG presents a risky prospect while choosing the private good ensures a guaranteed outcome. A risk averse subject is therefore expected to be is less inclined to contribute to the PG than a risk neutral subject. To avoid this bias, our binary choice formulation allows us to introduce risk into the private good, thereby reducing the potential for risk aversion to affect contribution behavior.

More precisely, let us denote the agents by $i \in \{1, 2, 3\}$. The common value of the PG is given by $P = s_1 + s_2 + s_3$ where each s_i is an independent draw from a uniform discrete distribution over the interval 0 to 100. Agent i observes only signal s_i. The value of the private good, V_i, differs for each agent, and is given by $V_i = D_0 + D_{1,i} + D_{2,i}$, where D_0 is exogenous, common, and common knowledge across all three agents. The six other signals, $D_{j,i}$ for $j \in \{1, 2\}$ and $i \in \{1, 2, 3\}$, are individual agents' private information and are each independent draws – also from a uniform discrete distribution over the interval 0 to 100. Therefore, after observing their own signals, each agent knows that the value of the PG is a known value plus two iid draws from a uniform distribution, and that the value of the private good is also a known value plus two iid draws from the same uniform distribution.

In our dynamic treatment, the values of both the PG and private goods are determined in exactly the same manner as the static treatment. The only difference is that decision making occurs in three stages, with simultaneous decisions within each stage. In the first stage, each agent has the option to select either the PG or private good. If an agent selects the PG, then the decision is final and is revealed to others in the group, and that agent does not participate in stages two or three.[5] Signals always remain private information. If an agent selects the private good in stage one, then in stage two they observe how many other group members selected the PG in stage one. In this second stage they may switch to select the PG or continue to choose the private good. Once

[5]Revealing prior commitments to the PG is analogous to the continuously updated cumulative prior contributions made by others on crowdfunding sites such as Kickstarter.

again, if the agent selects the PG then the decision is final and the agent does not participate in stage three. Agents who selected the private good in both stages one and two observe the number of PG decisions made by others in stages one and two and then, for the third and final time, they select either the PG or private good.

2.1 Equilibrium and hypotheses

Multiple equilibria exist in all of our treatments. For example, given the requirement that at least two agents must select the PG for it to be provided, it is always an equilibrium for no agent to select the PG. Previous research has identified conditions in which this type of inefficient equilibrium is not trembling hand perfect in private value environments (Bagnoli and Lipman, 1989). We focus on (symmetric) equilibria in which the PG is provided with positive probability.

Appendix B presents details for the static and dynamic treatment equilibria for the experimental environment. Here we provide a short intuitive summary. For the symmetric distribution of signals, which we simplify to the uniform distribution over the interval 0 to 100 to facilitate subjects' understanding, a subject who chooses the private good will earn, in expectation $\mathbb{E}[V_i] = \mathbb{E}[D_0 + D_{1,i} + D_{2,i}] = D_0 + 100$. A subject who ignores selection effects and chooses the PG would expect to earn $\mathbb{E}[P] = \mathbb{E}[s_1 + s_2 + s_3] = s_1 + 100$, if they ignore the fact that other agents choice of the PG is good news indicating the PG has higher value. This comparison between the private good and PG naive expected value suggests the simple but incorrect decision rule of selecting the PG if and only if $s_1 \geq D_0$. This is exactly the decision rule implied by cursed equilibrium (Eyster and Rabin, 2005), which is characterized by every agent making an optimal decision under the erroneous assumption that other players make decisions that are not conditioned on their private information.

Signal cutoffs for selecting the PG differ for sophisticated subjects who correctly account for the fact that other players select the PG when they have a higher signal. Consider first the static treatment. Given that the expected value of the PG is strictly increasing in player i's signal s_i, and that the value of the private good is independent of s_i, it is optimal for agents to use a cutoff strategy to choose the PG for signals above a particular threshold. This cutoff point is the signal where the agent is indifferent between the private good and the PG because they provide equal expected value. Whenever this cutoff point is positive, the PG choice of other agents is informative of their private signal, and changes its expected value. If the equilibrium cutoff is X, for example, then an outside observer expects that the average signal for agents who select the PG is $(X+100)/2$.

This exceeds the unconditional expected value of 50 for any $X > 0$. Consequently, when an agent's PG choice is pivotal (because at least one other agent also chose the PG) it has an expected value that exceeds the unconditional average. Agents should therefore choose the PG more frequently when they account for this selection. In other words, the selection effect lowers the threshold cutoff value for choosing the PG.

We show in Appendix B how much lower these Nash equilibrium cutoffs are lower than the cursed equilibrium cutoffs for any $D_0 > 0$. As shown in Table 2 for the three values of D_0 used in the experiment, the PG is chosen more often when agents correctly condition on the "good news" that they are more likely to be pivotal to obtain the PG when other agents have high signals and also opt for the PG.

Calculations are more complex for the dynamic treatment, because knowledge that other agents did or did not choose the PG in previous stages affects the estimates of the PG value. Appendix B provides derivations for the subgame perfect Nash equilibrium using backward induction, where agents with sufficiently high signals select the PG in early stages rather than delaying.[6] Two forces determine the first stage cutoff value. First, there is an option value from deferring a decision to select the PG: the longer I wait, the more I can infer about the private signals of others. The option value of waiting pushes the first stage cutoff value higher in the dynamic treatment (relative to the static treatment).

Second, there is a signaling effect: if I have a good signal I wish to communicate this to others, and induce their entry, by entering as soon as possible. The signaling effect increases the value of selecting the PG immediately for high private signals (as this will encourage entry by others and increase the chances that the PG is provisioned) but decreases the value of selecting the PG immediately for low private signals (as encouraging entry by others in this case can lead to inefficient provisioning of the PG).

Understanding the option value of waiting requires a subject to understand that there is information that can be extracted, in the future, from current decisions of other players. In contrast to equilibrium reasoning in the simultaneous treatment, the extraction of this information in the second stage does not require hypothetical thinking. On the other hand, the signaling effect requires hypothetical thinking about the future behavior of other players, but does not require an ability to extract information from a signal.

[6] Although equilibria exist with delay, they lead to lower expected payoffs and our experimental data provide no evidence consistent with them.

Private good base value (D_0):	0	30	70
Cursed equilibrium cutoff	0	30	70
Static equilibrium cutoff	0	25.0	52.1
Dynamic equilibrium cutoffs:			
Stage 1	47.7	58.9	73.0
1 Previous in Stage 1	17.5	33.7	55.7
1 Each in Stages 1 & 2	0	4.3	19.2
2 Previous in Stage 1	0	0	0
Public Good frequency:			
Cursed equilibrium	1.00	0.785	0.215
Static equilibrium	1.00	0.844	0.467
Dynamic equilibrium	0.843	0.655	0.361
Loss frequency (PG value < private good value):			
Cursed equilibrium	0.189	0.220	0.073
Static equilibrium	0.189	0.252	0.179
Dynamic equilibrium	0.096	0.148	0.112

Table 2: Equilibrium cutoffs and performance for all treatments.

To illustrate the signaling and information extraction effects, consider an example for the $D_0 = 30$ treatment shown in the middle column of Table 2. The cutoff to choose the PG for the static treatment is 25 rather than 30 in the cursed equilibrium, due to the positive selection discussed earlier, so a subject with a signal of 28, say, should choose the public good. For the dynamic treatment this subject with a signal of 28 should not choose the PG immediately. In equilibrium only agents with signals above 58.9 would choose the PG in stage 1, and if two group members do this then the third member can conclude that their expected signals are $(59 + 100)/2 = 79.5$. The updated expected value for the PG when the third member has a signal of Y is therefore $Y + 79.5 + 79.5 = Y + 159$. This exceeds the expected value of the private good (which is always 130 in this treatment) for any Y, which is why the cutoff is 0 for this treatment when two agents choose the PG in stage 1. In general, as illustrated in Table 2, cutoffs decline as more group members choose the PG in earlier stages. Due to the greater information dissemination from the sequential PG decisions, in the no-delay equilibrium players choose the PG more often in the rounds where it is efficient to do so in the dynamic treatment relative to the static treatment.

In addition to the equilibrium cutoffs for the static and dynamic treatments, Table 2 also summarizes the likelihood of PG choices in the static Nash, dynamic SPNE and cursed equilibrium, and expected frequency of inefficient PG choices (due to lower earnings than the private good) based on the uniform distribution of signal draws. These treatment differences lead to the following hypotheses that are based on proper contingent reasoning and equilibrium choices, as well as naive (unconditional) cursed beliefs for Hypotheses 1 and 3.

Hypothesis 1: (Outcomes) (a) The frequency of selecting the PG decreases as the private good base value (D_0) increases; and (b) the probability of a subject (inefficiently) receiving the PG when the private good has a higher value is larger for $D_0 = 30$ than either $D_0 = 0$ or $D_0 = 70$.

Hypothesis 2: (Outcomes) (a) The PG is chosen more frequently in the static than the dynamic treatment; and (b) the PG is chosen when it has a lower value than the private good more frequently in the static than the dynamic treatment.

Hypothesis 3: (cursed equilibrium) Estimated PG choice signal cutoffs correspond to the private good base value (D_0) for both the static and dynamic treatments.

The remaining hypotheses are based on Nash rather than cursed equilibrium.

Hypothesis 4: (a) Subjects choose the PG with lower frequency in stage 1 of the dynamic treatment than in the static treatment; and (b) estimated signal cutoffs for choosing the PG are higher in stage 1 of the dynamic treatment than in the static treatment.

Hypothesis 5: (a) Subjects choose the PG at higher rates in later stages of the dynamic treatment if more other agents in their group have previously selected the PG; and (b) estimated signal cutoffs in the dynamic treatment decrease for later stages when more other agents in the group previously selected the PG.

Note that in the last two hypotheses each case part (a) of the hypothesis is closely related to part (b) of the hypothesis in the sense that, assuming subjects are using cutoff strategies, analysis part (a) holds if and only if part (b) holds in the limit as the number of observations per subject increases. We test both parts, however, for the following reasons. First, in a finite sample, variation in the signal draws may mean that part (a) is rejected even when subjects are using the theoretically predicted cutoffs. Second, if decision noise or data sparsity near the cutoff point leads to a misestimation of the cutoff rules used by subjects then part (b) may be rejected even when subjects are using an appropriate cutoff rule, and part (a) then serves as a useful double check.

2.2 Laboratory procedures

The experimental design varied the common, baseline value of the public good at three levels, $D_0 \in \{0, 30, 70\}$, and whether the binary PG choice was static or dynamic. The D_0 value varied between subjects, as it was kept constant at one of the three values throughout each experimental session. Each session included 20 consecutive rounds of the static treatment and 40 consecutive rounds of the dynamic treatment, in two blocks; the ordering was varied so exactly one half of the sessions in each D_0 treatment began with the dynamic treatment and one half began with the static treatment. The static and dynamic treatments were varied within subject in order to determine whether the types of (cursed) contingent reasoning failures are corellated across settings. Independent signals s_i and $D_{j,i}$ were drawn each round. We conducted twice as many rounds for the dynamic treatment in order to obtain a greater number of observations for stage two and three decisions in different subgames (0, 1 or 2 earlier PG choices by others). Conducting the dynamic treatment using the strategy method was not an option given our research objective to study contigent reasoning.

We collected data from a total of 144 subjects, with 48 subjects in each D_0 treatment. Subjects were randomly reassigned to new groups of 3 each round, out of matching groups of size 12, so each treatment included 4 independent observations. The subjects were all undergraduate students at Purdue University, recruited from a database of approximately 3,000 volunteers drawn across a wide range of academic disciplines and randomly allocated to treatment conditions using ORSEE (Greiner, 2015). The oTree program (Chen et al., 2016) was used for the implementation of the experiment. We used neutral framing in the experiment, referring to choices between the "Group Project" or the "Private Project." Details are provided in the written instructions given to subjects (see Appendix A).

These written instructions were read aloud at the start of the session by an experimenter, while subjects followed along on their own hardcopy. New complete instructions were distributed at the treatment switch (from simultanous to dynamic or vice versa), but only the changes were highlighted and read aloud. Each session also concluded with two short "acquiring-a-company" game choices (both paid) for a separate measure of subjects' contingent reasoning. Sessions lasted about 1 hour on average, including instructions and payment time. At the conclusion of each session earnings were paid out privately in cash at a pre-announced conversion rate from Experimental Dollars earned for one randomly-drawn round for the main PG provision task. Subjects earned \$26.69 on average per person, with an interquartile range of [\$21.68, \$28.21].

3 Results

3.1 Public Good Choices and Provision

Hypothesis 4(a) states that agents should choose the PG with lower frequency in stage 1 of the dynamic treatment than in the static treatment. Figure 1 summarizes these individual choices for different PG value signals, providing support for this prediction for all three D_0 treatments. The figure also documents that subjects choose the PG more frequently for the treatments with a lower base value (D_0), consistent with Hypothesis 1(a), and they choose the PG with low s_i signals at substantial rates only for the lowest $D_0 = 0$. Aggregate PG choices do not, however, exhibit the sharp shift at equilibrium threshold signal levels (indicated on the figure as vertical lines). We consider individual threshold strategies in Section 3.3.

For signal values below the static Nash equilibrium threshold ($s_i = 0$, $s_i = 25$ and $s_i = 52$ for $D_0 = 0$, $D_0 = 30$ and $D_0 = 70$, respectively) the Nash equilibrium prediction coincides for both the static and stage 1 dynamic treatments. Similarly, for signal values above the stage 1 dynamic Nash equilibrium threshold ($s_i = 48$, $s_i = 59$ and $s_i = 73$ for $D_0 = 0$, $D_0 = 30$ and $D_0 = 70$, respectively), the Nash equilibrium predictions also coincide for both treatments. Therefore, we should expect to see treatment differences only between these ranges; for signal values such that a subject in the static treatment should select the PG and a subject in the first stage of the dynamic treatment should select the private good. Figure 1 demonstrates that this is indeed the case, as the treatment differences are substantial for signals that fall between the two equilibrium cutoffs, denoted by vertical red lines, for all three values of D_0. Differences in PG choice frequencies for the static and dynamic treatments in these key signal ranges are highly statistically significant, based on linear probability models with standard errors clustered on individual subjects, controlling for time trends and treatment ordering. (Estimated $p - values < 0.01$ for all comparisons.)

Result 1: Subjects choose the PG more frequently when the private good has a lower value (support for Hypothesis 1(a)) and more frequently in the static treatment than in stage 1 of the dynamic treatment (support for Hypothesis 4(a)).

Hypothesis 5(a) concerns later stage choices in the dynamic treatment, in particular that agents will choose the PG at higher rates in later stages of the dynamic treatment when more other agents in their group previously selected the PG. Table 3 reports a series of linear probability models of subjects' choice of the PG in the second stage, conditional on the number of others in the group who selected the PG in the first stage. The omitted case is for zero other group members

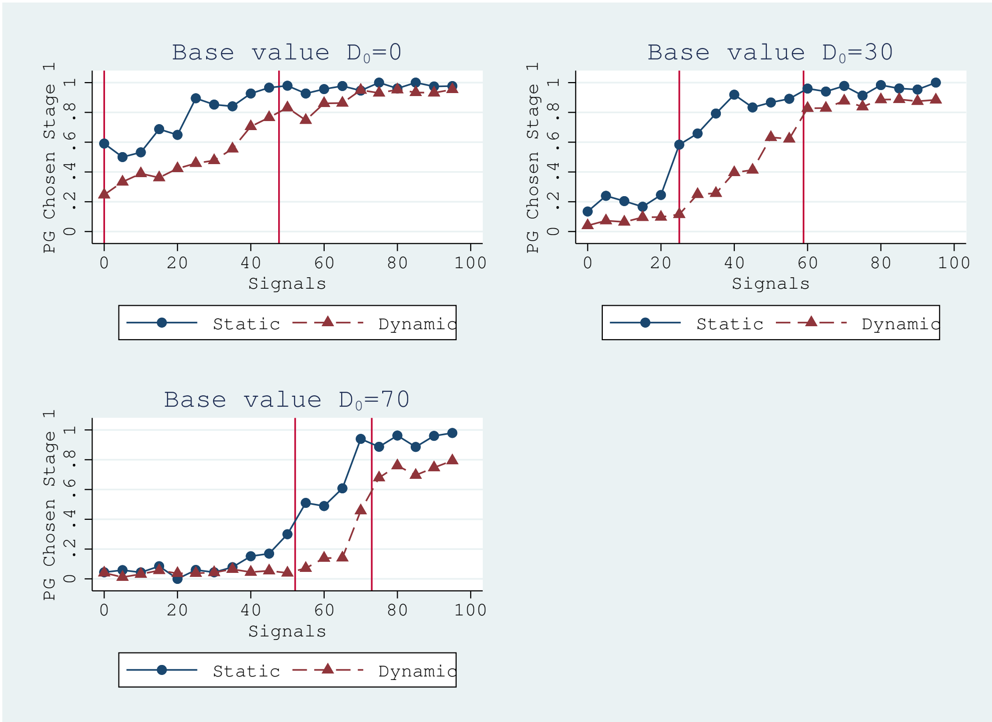

Figure 1: PG Choice Frequency by Signal, Static and Stage 1 Dynamic Treatments

selecting the PG in stage 1. The models include the own received signal (s_i) to control for the nonrandom selection (lower signal draws) for subjects to reach the later stages without having previously committed to the PG, and they also control for a time trend and treatment ordering.

The odd numbered columns report estimates without additional controls, while the even numbered columns add demographic characteristics as well as responses on the "acquiring-a-company" questions asked of subjects at the end of their session. We employ multiple elicitations of this separate measure of individuals' comprehension of contingent reasoning and apply the *obviously related instrumental variables* method of Gillen et al. (2019) to attenuate measurement error. Results are similar with and without these controls.

The regression results show that having two rather than just one other subject choosing the PG previously has a particularly strong impact on the stage 2 PG decisions. For all six models the coefficient on two previous entries is significantly greater than for one previous entry (all p-values< 0.001). Subjects with higher signals are also significantly more likely to choose the PG in stage 2.[7] To summarize:

[7]Similar results obtain for stage 3 decisions, although we do not include them in Table 3 because the number of observations is lower and so the statistical significance is weaker, and the selection effect of nonrandom, low signal choices in the third stage is much stronger.

	Stage 2 PG for $D_0 = 0$		Stage 2 PG for $D_0 = 30$		Stage 2 PG for $D_0 = 70$	
	(1)	(2)	(3)	(4)	(5)	(6)
One other previous PG choice	0.083 (0.044)	0.054 (0.114)	0.063 (0.034)	0.064* (0.032)	0.135*** (0.021)	0.134*** (0.020)
Two other previous PG choices	0.614*** (0.057)	0.603*** (0.096)	0.549*** (0.045)	0.559*** (0.044)	0.609*** (0.054)	0.602*** (0.053)
Own signal (s_i)	0.0045*** (0.0011)	0.0049** (0.0018)	0.0076*** (0.0016)	0.0075*** (0.0014)	0.0062*** (0.0007)	0.00611*** (0.0007)
Round number t in treatment	0.0022 (0.0016)	0.0020 (0.0023)	0.0006 (0.0011)	0.0009 (0.0011)	-0.0001 (0.0007)	-0.0002 (0.0007)
Treatment order in session	0.0025 (0.0653)	0.0607 (0.1980)	-0.0359 (0.0426)	-0.0692* (0.0329)	-0.0158 (0.0238)	-0.0111 (0.0243)
Demographic and ATC game controls	No	Yes	No	Yes	No	Yes
N	604	604	968	968	1439	1439
adj. R^2	0.347		0.343		0.334	

Standard errors (clustered on individual subjects) in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Stage 2 Public Good Choices in Dynamic Treatment.

Result 2: Subjects choose the PG at higher rates in later stages of the dynamic treatment if more members of their group have previously chosen the PG, particularly for two previous PG choices (support for Hypothesis 5(a)).

Hypothesis 2 concerns the overall provision rate for the PG, an outcome which depends on the decisions of multiple agents given that at least two people must select the PG for it to be provided. Table 4 displays the rate of PG provision by both treatment type (dynamic or static) and the value of D_0. No discernable time trend exists for the static treatment, but the dynamic treatment exhibits a modest decreasing time trend as indicated in the disagregated early (periods 1-20) and later (periods 21-40) shown in the table. Hypothesis 2(a), that the PG is provided at a higher rate in the static than the dynamic treatment, is rejected by the data. In opposition to the hypothesis, the rate of PG provision is actually higher in the dynamic than the static treatment for the $D_0 = 30$ and $D_0 = 70$ treatments.[8] In terms of point predictions, the average PG provision rate

[8]We establish this using a linear probability model with clustered standard errors, controlling for time trends and treatment ordering. The differences between the static and dynamic treatment are significant at the 1-percent level for both the $D_0 = 30$ and 70 treatments, either considering only the comparable first 20 periods of each treatment

Private good base value (D_0):	0	30	70
Public Good frequency:			
Cursed equilibrium	1.00	0.785	0.215
Static equilibrium	1.00	0.844	0.467
Dynamic equilibrium	0.843	0.655	0.361
Static actual (all periods 1-20)	0.835	0.645	0.283
(standard error of mean)	(0.012)	(0.015)	(0.015)
Dynamic actual (all periods 1-40)	0.836	0.672	0.359
(standard error of mean)	(0.008)	(0.011)	(0.011)
Dynamic actual (periods 1-20)	0.857	0.728	0.373
(standard error of mean)	(0.011)	(0.014)	(0.016)
Dynamic actual (periods 21-40)	0.816	0.617	0.346
(standard error of mean)	(0.013)	(0.016)	(0.015)
Loss frequency (PG value $<$ private good value):			
Cursed equilibrium	0.189	0.220	0.073
Static equilibrium	0.189	0.252	0.179
Dynamic equilibrium	0.096	0.148	0.112
Static actual (all periods 1-20)	0.152	0.165	0.094
(standard error of mean)	(0.012)	(0.012)	(0.009)
Dynamic actual (all periods 1-40)	0.158	0.178	0.123
(standard error of mean)	(0.008)	(0.009)	(0.008)
Dynamic actual (periods 1-20)	0.169	0.188	0.125
(standard error of mean)	(0.012)	(0.013)	(0.011)
Dynamic actual (periods 21-40)	0.147	0.169	0.122
(standard error of mean)	(0.011)	(0.012)	(0.011)

Table 4: Equilibrium and realized PG provision and overprovision for all treatments.

is remarkably close to the Nash Equilibrium prediction for the dynamic treatment. Surprisingly, the PG provision rate for the static treatment is also similar to the dynamic equilibrium prediction, except for $D_0 = 70$ where the rate is even lower. The downward bias in the static treatment for the PG provision rate is in the direction of the cursed equilibrium for $D_0 > 0$, particularly for $D_0 = 70$.

Hypothesis 2(b) concerns errors in PG provision, in particular the provision of the PG to agents in rounds where, ex-post, the private good would have generated a higher payoff. One of the key theoretical advantages of the dynamic mechanism is that, in equilibrium, such over-provision of the PG is reduced compared to the static mechanism. The lower half of Table 4 shows the overall

or all periods.

average over-provision. For the $D_0 = 30$ and 70 treatments the error rates are similar to the dynamic equilibrium rates in both the static and dynamic treatments. Differences between static and dynamic treatment error rates are not statistically significant in the $D_0 = 0$ and 30 treatments, but for the $D_0 = 70$ treatment the over-provision errors are significantly *greater* in the dynamic than static treatment.[9] Thus, the data indicate that the symmetric Nash equilibrium does not accurately predict the rates of provision and over-provision of the PG in aggregate, particularly for the static treatment.

Hypothesis 1(b) states that errors in PG provision are greater in the $D_0 = 30$ treatment than either the $D_0 = 0$ or $D_0 = 70$ treatments. Table 4 indicates that mean error rates are consistent with this hypothesized ordering, and regression-based statistical tests indicate that these differences are statistically significant for the $D_0 = 30$ to $D_0 = 70$ comparison in both the static and dynamic treatments (p-values< 0.001 in both cases). For the $D_0 = 30$ to $D_0 = 0$ comparison the $D_0 = 0$ treatment has significantly more errors in the dynamic treatment (one-tailed p-value$= 0.05$) but no significant difference in the static treatment. Thus, this hypothesis receives support in 3 out of the 4 relevant comparisons.

Result 3: The PG provision rate in the dynamic treatment is greater than or equal to the rate in the static treatment, and the over-provision (error rate) in not lower in the dynamic treatment (Hypothesis 2 not supported). Consistent with Hypothesis 1(b), over-provision of the PG is greatest in the intermediate ($D_0 = 30$) treatment.

Overall, the Nash equilibrium provides a useful approximation for aggregate behavior in the dynamic treatment but not for the static environment. Nevertheless, aggregated outcomes may mask the choice of strategies at the individual level. In particular, studying the strategy choices of subjects may help explain the bias away from Nash outcomes, and towards outcomes predicted by Cursed equilibrium, in the static treatment. Thus, we now turn to a study of the strategies used by subjects in both our static and dynamic treatments.

3.2 Estimating strategies: Cutoff Intervals

Recall that a rational agent will use a cutoff rule in all scenarios: if the agent observes a signal s_i for the PG that is above some threshold the agent will prefer the PG and otherwise prefers the private

[9]This conclusion is based on the same type of regression summarized in the previous footnote. Differences are significant at the two-percent level for $D_0 = 70$ regardless of whether all periods or only the first 20 periods are compared.

good. Therefore, we summarize each subject by four points x_S, x_D, x_1 and x_2, where x_S denotes the cutoff in the static treatment, x_D denotes the cutoff in the first round of the dynamic treatment, x_1 denotes the cutoff in the dynamic treatment when observing that one other player has already committed to the PG, and x_2 denotes the cutoff in the dynamic treatment when observing that both other players have committed to the PG.[10]

The subjects' binary choice data do not reveal their cutoff points directly. We therefore infer their cutoffs using a deterministic process that provides interval identification of each cutoff point. Our procedure is maximally efficient: assuming a subject is using a cutoff rule we identify, conditional on the observed data, the smallest possible interval that contains the cutoff point. The procedure is as follows, for a subject who observes k signals. Order the signals observed by the subject from smallest to largest, labeled s^1 through s^k, and denote the ordered set by S. The data are summarized by the mapping $d : S \to \{0, 1\}$ where $d(s^k) = 1$ indicates that the subject selected the PG when signal s^k was observed, and $d(s^k) = 0$ indicates that the subject selected the private good. Next, identify the set of $k + 1$ possible intervals $I = \{I^0 = [0, s^1], I^1 = [s^1, s^2], \ldots, I^k = [s^k, 100]\}$. Then, for each interval, calculate an error index $E(I^j)$ for $j \in [0, k]$ as follows:

$$E(I^j) = \begin{cases} \sum_{i=j+1}^{k} 1 - d(s^i) & \text{if } j = 0 \\ \sum_{i=1}^{j} d(s^i) + \sum_{i=j+1}^{k} 1 - d(s^i) & \text{if } 1 \leq j < k \\ \sum_{i=1}^{k} d(s^i) & \text{if } j = k \end{cases}$$

Finally, if $\arg\min_{j\in[0,k]} E(I^j)$ is unique, then we conclude that the cutoff point must lie in I^j. If $\arg\min_{j\in[0,k]} E(I^j)$ is not unique, then we conclude that the cutoff point must lie in the $[s_{\underline{j}}, s_{\overline{j}+1}]$ where $\underline{j}$ and $\overline{j}$ are the smallest and largest minimizers, and we adopt the convention that $s_0 = 0$ and $s^{k+1} = 100$.

We illustrate the operation of this procedure by highlighting some examples from our data in Figure 2. Each of the examples below are taken from the static $D_0 = 30$ treatment, where each subject was observed making 20 decisions.

Consider subjects 2 and 37. The behavior for each of these subjects is perfectly consistent with the use of a cutoff strategy, and the error index for each subject obtains a minimum value of 0. For

[10] Theory suggests that subjects might use a different cutoff in the cases where both other players are observed to commit to the PG in the first stage, and where one player commits in the first stage and another player commits in the second stage (cf Table 2). We do not have enough observations per subject to observe stage 3 cutoffs reliably in the data, so we instead focus only stage 2 decisions. This does not affect Hypothesis 5.

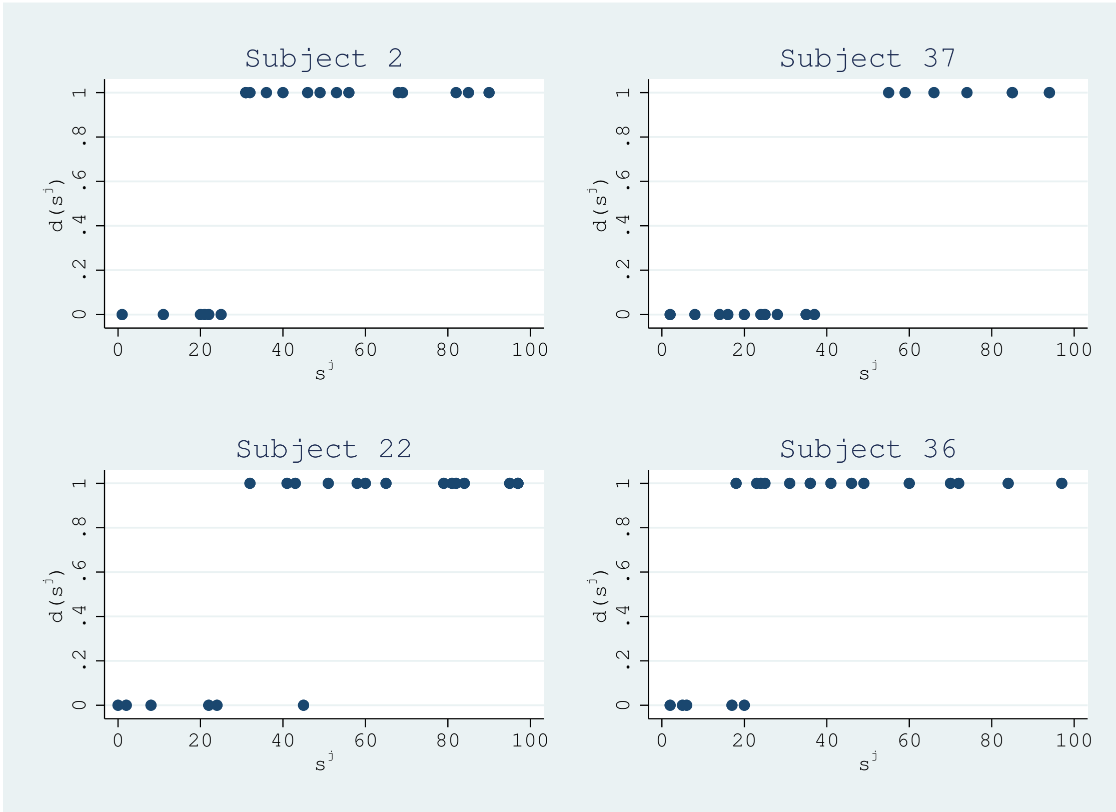

Figure 2: Individual PG Choices for Example Subjects in the Static Treatment

subject 2 the interval that minimizes the error index is reasonably tight, with this subject using a cutoff that lies between 25 (the largest signal at which the private good was chosen) and 31 (the smallest signal at which the PG was chosen). Subject 37 exhibits a wider interval, and is using a cutoff that lies between 37 and 55. The precision with which we can infer the cutoff rule of a subject varies across subjects, and depends on the subject's behavior as well as the random draw of signals that the subject received.

Subject 22 also appears to be using a cutoff rule, but may have made a "tremble" or mistake in one round. The error index for this subject obtains a minimum value of 1 over the interval $[24, 32]$. Subject 36 also has a minimum error index of 1 but, unlike subject 22, the interval that minimizes the error index is not unique. The error index for subject 36 obtains the value of 1 over the interval $[17, 18]$ and the interval $[20, 23]$. We therefore estimate the cutoff for this subject as belonging to an interval that spans from 17 to 23.

On aggregate, our subjects do appear to be implementing cutoff rules. In the static treatment, where each subject is observed making 20 decisions, the median error index is 0.5 and the mean is 1.1. In the first period of the dynamic treatment, where each subject is observed 40 times, the median error index is 2 and the mean is 3.0. The larger error indices in the dynamic treatment are a function of the larger number of decisions leading to more opportunities to make a "mistake." If we restrict attention to the final 20 rounds of the dynamic treatment, the median error index drops to 0 and the mean drops to 0.8. In the later stages of the dynamic treatment the number of observations per subject is smaller, given that many subjects commit to the PG in the first period. Across all 40 rounds of the dynamic treatment, we observe each subject making, on average, 9.1 decisions in which one other subject has already committed to the PG and 5.0 decisions in which both other subjects have already committed. In both of these cases the error indices are expectedly small, with a median of 0 in both cases and mean error indices of 0.6 and 0.4, respectively.

Finally, we adjust the observed intervals in the dynamic treatment to be consistent with an intuitive monotonicity condition: the cutoff at which a subject selects the PG should be non-increasing across stages of the dynamic treatment. We illustrate the problem with a simple, and typical, example. Suppose that, in the dynamic treatment, a subject is using a decision rule such that she chooses the PG whenever $s_i >= 30$ in the first stage and that, given her observed signals, the algorithm above assigns her an interval of $[28, 35]$. Then, in the second stage, the subject once again uses a decision rule such that she chooses the PG whenever $s_i >= 30$. This implies that we never observe the subject select the PG in the second stage, and the algorithm assigns an interval of $[28, 100]$. After applying our monotonicity condition, we adjust the second stage interval to $[28, 35]$. In some cases, subject behavior is not compatible with monotonicity. This usually occurs for subjects who have a high error index, suggesting that these subjects are not playing a cutoff strategy in the first place.

In the analysis that follows we exclude subjects who violate the monotonicity conditions, and subjects who have a total error index of 6 or more.[11] Out of 144 subjects, 23 are excluded for monotonicity violations and a further 14 are excluded for an elevated error rate.

3.3 Analysis of strategies

We use interval regression techniques to estimate average cutoff strategies for the various stages and treatments. The true, unobserved, cutoffs for each subject in each treatment scenario are

[11] We calculate the total error index by aggregating errors made in the static treatment, all rounds of the dynamic treatment where either one or two opponents have been observed to select the PG, and the final 20 rounds of the first stage of the dynamic treatment.

modeled as a normally distributed latent variable.[12] The interval regression estimates maximize the likelihood that the unobserved cutoffs lie within the intervals calculated above.

Once again we estimate our model both with and without demographic controls, including performance in the acquire a company game. Table 5 presents the average predicted cutoff value for each treatment. We estimate the model separately for $D_0 = 0, 30$ and 70. The estimated coefficients are remarkably robust to the demographic controls as are the standard errors, which are clustered at the subject level.

The stars in Table 5 denote the statistical significance of the p-value testing the null hypothesis that the relevant estimated cutoff is equal to the estimated cutoff of the first stage of the Dynamic treatment. It is clear from the table that the average cutoff in the first stage of the Dynamic treatment is greater than the average cutoff in the static treatment, and also greater than the cutoff in the Dynamic treatment after others are observed to select the PG.

Result 4: Estimated signal cutoffs for choosing the PG are higher in stage 1 of the dynamic treatment than in the static treatment (support for Hypothesis 4(b)) and estimated signal cutoffs decrease for later stages in the dynamic treatment when more other agents in the group choose the PG (support for Hypothesis 5(b)).

The largest deviations from the Nash equilibrium point predictions occur in the Dynamic treatment when both opponents have already chosen the PG. The NE predicts that the PG should always be selected in this case, i.e., a cutoff of 0, but Table 5 shows that the estimated cutoff values are substantially larger than 0 in all treatments ($p < 0.001$ for all cases.) Subjects are under-reacting to the information encoded in observing both opponents selecting the PG. Given the observed behavior in the first stage, the best response cutoff is 1, 0 and 0 in the $D_0 = 0, 30$ and 70 treatments respectively.[13]

Subjects are much closer to best responding to the information encoded in observing only one opponent choose the PG, although there is still some under-reaction in the $D_0 = 30$ and 70 treatments. Using the estimated cutoff in the first stage of the Dynamic treatment, and the best response condition Equation 4, we calculate the best response cutoff after observing one selection of the PG to be 24, 35 and 54 in the $D_0 = 0, 30$ and 70 treatments, respectively. The estimated

[12]The treatment scenarios are the Static treatment, the first stage of the Dynamic treatment, the Dynamic treatment after observing one other subject select the PG, and the Dynamic treatment after observing both other subjects select the PG. For brevity, we will use the term "treatments" to refer to these four scenarios, even though formally they are not all distinct treatments.

[13]$p < 0.001$ for all treatments. Best responses are calculated using Equation 2 of Appendix B.

	$D_0 = 0$		$D_0 = 30$		$D_0 = 70$	
	(1)	(2)	(3)	(4)	(5)	(6)
Static	13.09***	13.07***	28.46***	28.53***	59.81***	59.82***
	(2.667)	(2.526)	(1.837)	(1.686)	(1.486)	(1.312)
Dynamic	28.58	28.55	55.35	55.35	76.97	77.05
(Stage 1)	(3.161)	(3.174)	(2.694)	(2.645)	(1.930)	(2.036)
Dynamic	24.30*	24.23*	42.35***	42.50***	59.58***	59.54***
(One previous entrant)	(3.016)	(2.941)	(2.095)	(2.074)	(1.462)	(1.479)
Dynamic	8.722***	8.746***	15.75***	15.49***	33.11***	33.13***
(Two previous entrants)	(1.806)	(1.596)	(2.020)	(1.854)	(4.321)	(4.094)
Treatment order controls	Yes	Yes	Yes	Yes	Yes	Yes
Demographic and ATC game controls	No	Yes	No	Yes	No	Yes
N	128	128	148	148	152	152

Standard errors, clustered at the subject level in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 5: Average predicted cutoff value by treatment and number of observed entrants, calculated via interval regression. p-values are calculated from a test of equality of the relevant cutoff with the Dynamic (Stage 1) cutoff. Treatment order controls included in all regressions, demographic and contingent reasoning (ATC) controls included in even-numbered columns. Contingent reasoning controls are included using the obviously related instrumental variables technique of Gillen et al. (2019). Sample restricted to subjects who do not violate the monotonicity constraints and exhibit 5 or fewer choice errors.

values of these cutoffs are 24, 42 and 60, respectively.[14]

Figure 3 plots the estimated cutoff values along with equilibrium predictions. Nash equilibrium predicts the direction of treatment effects across variation in both D_0 and the timing of the game. As a point prediction Cursed equilibrium does not perform well in our data, and Cursed equilibrium Hypothesis 3 is rejected in 11 out of 12 cases. However, Cursed equilibrium does help to organize our data in the static treatment, where deviations from Nash equilibrium are in the direction of Cursed equilibrium (whenever the Nash and Cursed equilibrium differ). In the case of $D_0 = 30$, the 95% confidence interval of the average cutoff value covers the Cursed equilibrium but not the Nash equilibrium. In the Dynamic treatment, the Nash equilibrium point predictions perform well in the first stage, across the $D_0 = 30$ and $D_0 = 70$ treatments, and after observing one other player select the PG but not otherwise.

[14] $p = 0.937$, $p < 0.001$ and $p < 0.001$, respectively, for tests comparing the estimated cutoffs to these best responses.

Figure 3: Estimated cutoff values with 95% confidence intervals, controlling for demographic characteristics. Cursed equilibrium prediction denoted by dashed red lines, and Nash equilibrium predictions denoted by solid green lines.

Figure 4 plots the CDF of the midpoint of the cutoff intervals described in Section 3.2 for the Static treatment and the first stage of the Dynamic treatment. There is a clear rightward shift in the distribution of cutoffs in the Dynamic treatment, relative to the Static treatment, along with a substantial amount of within treatment heterogeneity.[15] Figure 5 focuses only on the Dynamic treatment, and demonstrates that the heterogeneity is not a result of learning across rounds. Figure 5 plots the CDF of the midpoint of the cutoff intervals for three cases: the entire 40 rounds of the Dynamic treatment, the first 20 rounds only, and the last 20 rounds only. As the figure makes clear, the heterogeneity of cutoff midpoints persists across time, with extremely little learning occurring in the $D_0 = 70$ treatment in particular.

To further investigate the source of the observed heterogeneity of strategies, in the Static treatment, where we do not have enough data to reliably split the sample into early and late rounds,

[15]The within-subject correlation of the dynamic and static treatment cutoff interval midpoints is weak, ranging between -0.08 and 0.10 depending on how the cutoffs are normalized and never significantly different from zero.

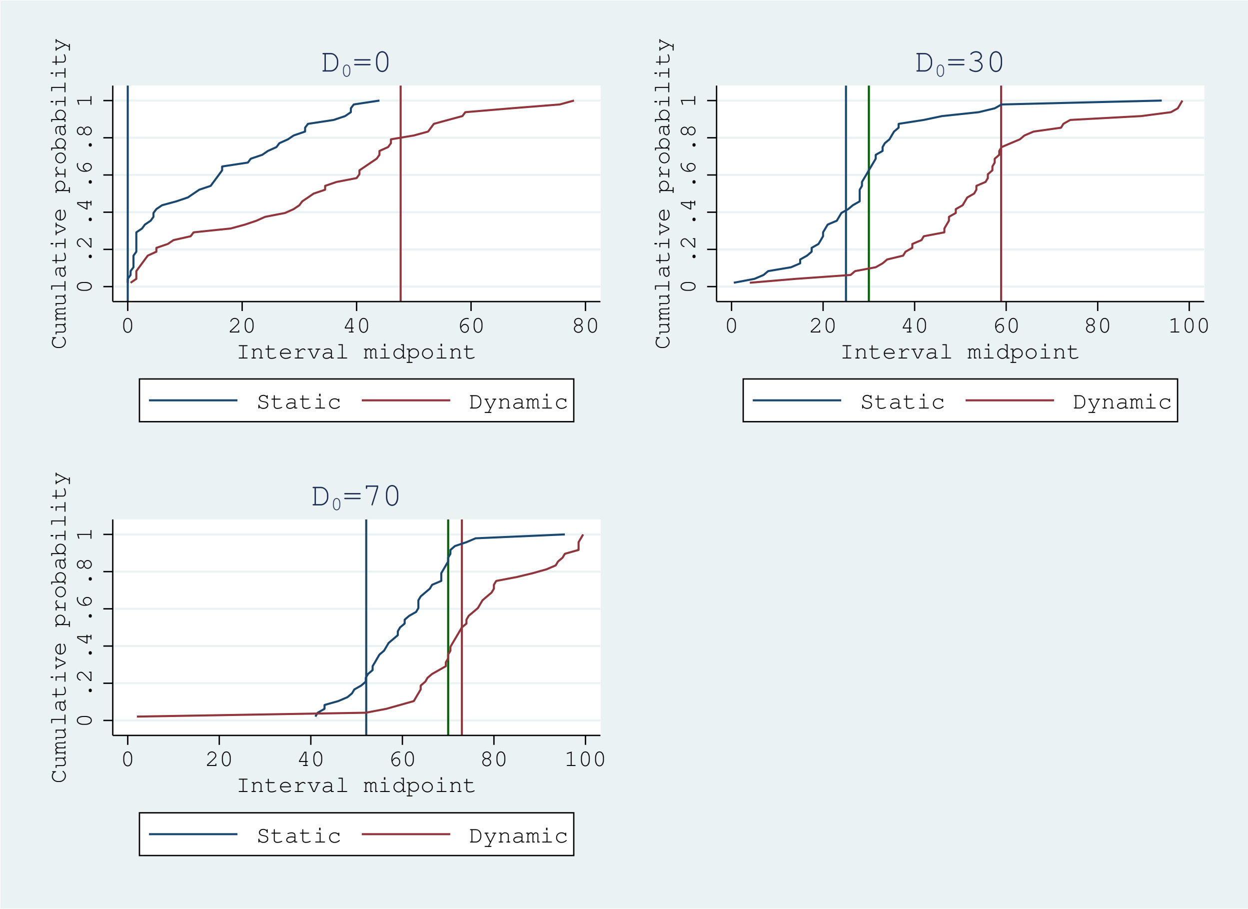

Figure 4: Cumulative Density Functions of the midpoint of the cutoff intervals (see Section 3.2). Vertical lines denote equilibrium predictions for Static Nash equilibrium, Dynamic Nash equilibrium and Cursed equilibrium in navy, maroon and dark green, respectively.

we calculated the empirical best response (EBR) cutoff value for each subject. For each subject we simulated 10 million rounds of the game where, in each round, the subject was matched against two out of the 11 cutoff strategies belonging to the 11 opponents in the subject's matching group.[16] Very little variation exists in these estimated EBR cutoffs calculated within each value of D_0. For the $D_0 = 0$ treatment the EBR cutoffs range from 0 to 2, with a mean of 0.1. For the $D_0 = 30$ treatment the EBR cutoffs range from 21 to 26, with a mean of 23.5, and for $D_0 = 70$ treatment the EBR cutoffs range from 45 to 50, with a mean of 47.6.

In both treatments where the Static Nash equilibrium is greater than 0, the mean EBR cutoff lies below the Nash equilibrium cutoff, while the mean observed cutoff is above the Nash cutoff. This

[16]The probability of a subject receiving a signal draw that induces different behavior between a cutoff strategy at x and $x+1$ is $\frac{1}{101}$. We therefore need a large number of rounds to reliably discriminate between cutoffs that are near optimal.

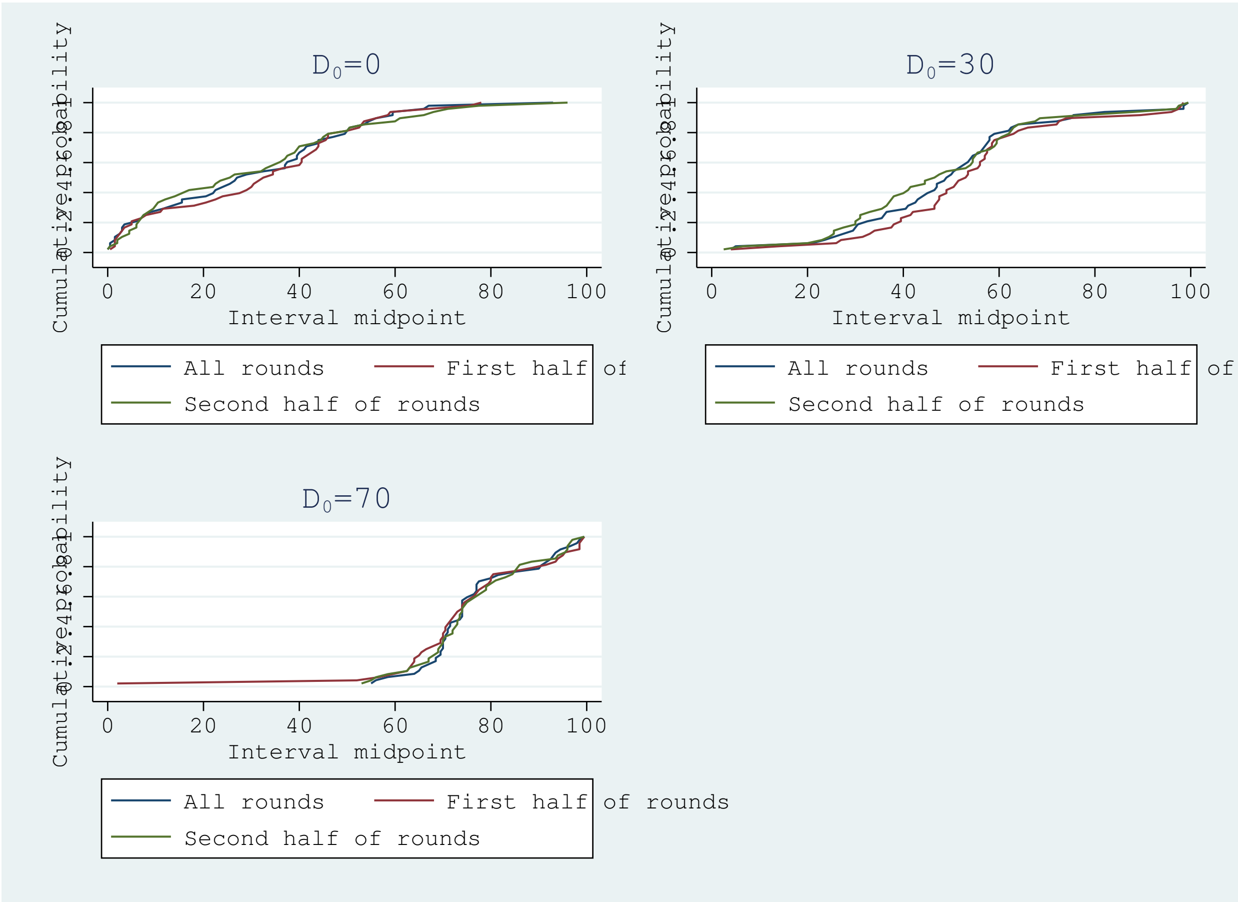

Figure 5: Cumulative Density Functions of the midpoint of the cutoff intervals (see Section 3.2) for the first stage of the Dynamic treatment. Intervals calculated for all rounds, rounds 1 to 20 and rounds 21 to 40.

reflects the underlying structure of the game, where cutoffs are strategic substitutes. If one player raises his cutoff then his opponents should lower their own cutoffs because, conditional on the player selecting the PG, the expected value of the PG is larger. An implication of this analysis is that, because most subjects are using cutoffs above the Nash cutoff, measuring the degree of cursedness relative to the EBR leads to higher measured levels of cursedness than measuring relative to the Nash equilibrium.

Thus, neither the observed bias away from Nash equilibrium, nor the substantial heterogeneity, in our data can be explained by subjects best responding to the empirical distribution they faced during the experiment. We therefore conclude that the systematic deviations from Nash equilibrium observed in the data are a consequence of the mental reasoning that subjects applied in this public goods game. The reasoning of subjects displays both substantial heterogeneity and a bias towards cursed thinking.

4 Conclusion

As noted in the introduction, it is rare for threshold PG environments to involve simultaneous decision making. Nevertheless, the incentives for a threshold PG game can be modified to reflect either our Dynamic or Static treatments by revealing, or not revealing, respectively, the current level of contributions in real time. The choice of information structure is, therefore, an important and easily manipulable policy variable for the designer of a threshold PG mechanism.

Our theoretical analysis suggests that there is an important tradeoff between the Static and Dynamic mechanisms. The Static mechanism generates a higher rate of PG provision, but achieves this, in part, by increasing the proportion of times where the PG is provisioned inefficiently. The Dynamic mechanism can improve PG provision choices, but introduces greater complexity. For a company such as Kickstarter this implies a tradeoff between revenue (which is a function of the number of projects that are financed) and long term reputation (which is harmed when consumers purchase a poor product).

Our experimental results do not support this theoretical tradeoff. The rate of PG provision in the Dynamic treatment is, if anything, slightly *higher* than in the static treatment and we do not find a difference in the rate of "mistakenly" provisioned PG. Further, because the threshold for commiting to the PG is substantially higher in the first stage of the Dynamic treatment than the Static treatment, there will be fewer near-misses (where a PG almost, but not quite, reaches the funding threshold) in the Dynamic mechanism. Each of these properties suggests that the Dynamic mechanism is likely to be more desirable from a practical standpoint. It is therefore perhaps no accident that Kickstarter and other crowdfunding sites update previous contribution continuously to promote information dissemination in their versions of a dynamic mechanism.

The experimental results provide clear support for the *complexity* explanation, and not the *unawareness* explanation, for the failure of contingent reasoning that has been widely observed in the literature. We find evidence, supporting the previous literature, for failures of contingent thinking: decisions in the Static treatment are biased towards Cursed equilibrium. We also find that strategies in the first stage of the Dynamic treatment differ significantly from strategies in the Static treatment. This suggests that the failures of contingent thinking are generated by the *complexity* of the contingent thinking problem, and not by *unawareness*. The stability of behavior across rounds throughout our experiment provides further evidence that the subject's awareness of the contingent thinking problem was not induced by participation in our experiment. Human beings are innately aware of the need for contingent thinking, and that they should take actions that allow their future selves to make use of valuable contingent information, but have difficulty in effectively solving hypothetical contingent thinking problems.

A Experiment Instructions

Part One **(Differences between Dynamic and Static highlighted in bold)**

Overview

This is an experiment in the economics of decision-making. The amount of money you earn depends partly on the decisions that you make and thus you should read the instructions carefully. The money you earn will be paid privately to you, in cash, at the end of the experiment. A research foundation has provided the funds for this study.

There are two parts to this experiment. These instructions pertain to Part 1A of the experiment. Once Part 1 is complete, the instructions for Part 2 will be distributed. Part 1 of the experiment is divided into many decision "periods." For Part 1, you will be paid your earnings in one, randomly selected, period. The period for which you will be paid shall be announced at the end of the experiment. Each decision you make is therefore important because it has a chance to affect the amount of money you earn.

In each decision period you will be grouped with two other people, who are sitting in this room, and the people who are grouped together will be randomly determined each period. You will be in a "matching group" of twelve people. You will only ever be matched with other people in the same "matching group" as yourself, which means that there are at most eleven other people you could be matched with each period.

You will make decisions privately, that is, without consulting other group members. Please do not attempt to communicate with other participants in the room during the experiment. If you have a question as we read through the instructions or any time during the experiment, raise your hand and an experimenter will come by to answer it.

Your earnings in Part 1 of the experiment are denominated in experimental dollars, which will be exchanged at a rate of 10 experimental dollars = 1 U.S. dollar at the end of the experiment.

Your Decisions

Part 1A of the experiment consists of 40 periods. **(20 periods for Static treatment)**

In each period, you will choose whether to receive earnings from the *group project* or you may instead choose to receive earnings from your *private project*. You will receive earnings *either* from the group or the private project, and *never* from both projects. Everyone in your group each period will make a similar decision. If you choose the group project, you only will receive earnings from the group project if at least one other person in your group chooses the group project. If you are the only one choosing the group project, then you receive earnings from the private project instead. The details of your earnings for these decisions are described below.

Group Project

In each period, a random number will be selected by the computer for you from a uniform distribution between 0 and 100. The uniform distribution means that the 101 possible values 0, 1, 2, ..., 99, 100 are equally likely. We will call this random number your signal. Each other member of your group will also get a signal randomly selected by the computer from this same distribution. We will call the signals of the three group members S1, S2 and S3. All signals are drawn *independently*, which means that no drawn signal can have any influence on any other signal draws. During each period, you will not observe the signals of the other members. Similarly, other members of the group will not observe any signal other than their own.

If you choose the group project, and if at least one other member of your group also chooses the group project, then you receive earnings that are equal to the sum of the signals of all three members of your group. We will call the sum of the signals of the three members of your group the value of the group project, or V:

V = S1 + S2 + S3

So, for example, if your signal is 50 and the other members of your group get signals of 25, and 86, then the sum of all three signals is:

V = 50 + 25 + 86 = 161

Thus, in this case, if you chose the group project and at least one other member of your group also choses the group project, then you would get 161 experimental dollars for that period.

If you choose the group project, but *no other* members of your group also choose the group project, then you receive earnings from your private project instead for that period. (In other words, if less than two of the three members of your group (including yourself) choose the group project, then all group members receive earnings from their private projects that period.) These private project earnings are described next.

Private Project

In each period, the baseline value of your private project is 70. This number is predetermined (i.e. not random) and is the same for all three members of your group. In each period, in addition, two random numbers will be selected by the computer for you from a uniform distribution between 0 and 100. We will call these two draws D2 and D3. These two numbers are drawn independently and will determine your earnings from the private project. (Other group members will receive their own random numbers, independently drawn, for their private projects.)

Like the group project, your private project value (P) comes from the sum of three values:

P = 70 + D2 + D3

You will only know the baseline value of 70 before you make your decision. So, for example, if the other two drawn values that you did not learn are D2=6 and D3=46, then your earnings from the project would be

P = 70 + 6 + 46 = 122

You will receive these private project earnings if either (1) you choose the private project or (2) you are the only person in your group who chooses the group project.

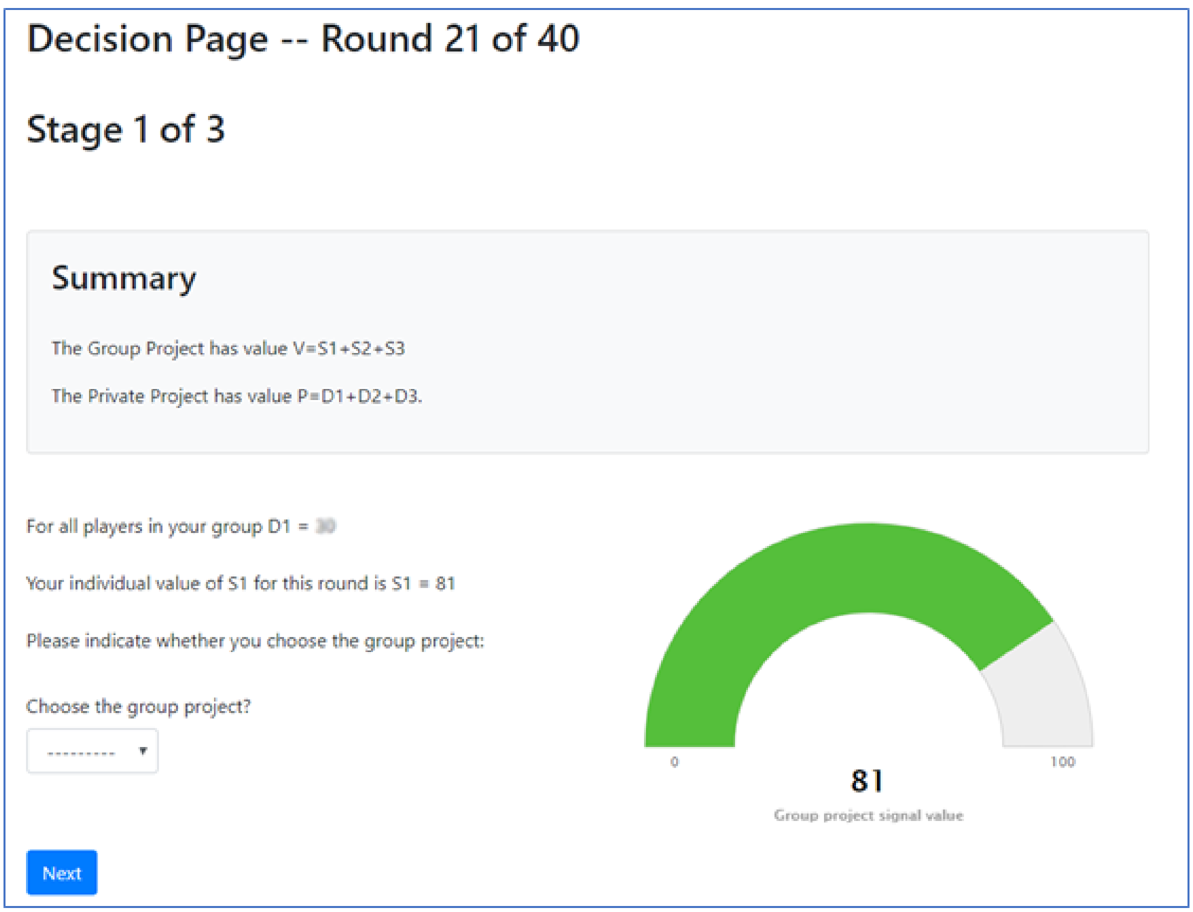

Figure 6: Decision Screen

Note that at the time of your choice, you only observe your own signal (S1, S2 or S3) of the group project value and the baseline number, 70, that determines part of the value of your private project. This is illustrated in your decision screen shown in Figure 6.

Three Choice Stages

You and other group members will have an opportunity to choose the group project in 3 sequential stages each period. If you choose the group project in an early stage you cannot switch to choose the private project instead in a later stage. But if you choose the private project in an early stage *you can switch* your choice to the group project in a later stage.

In Stage 1, everyone will make a first choice between the group and private project before learning the decisions of other group members.

In Stage 2, everyone will learn how many group members chose the group project in Stage 1, and those who have not yet chosen the group project may then switch to the group project. This is illustrated in Figure 7.

In Stage 3, everyone will learn how many group members chose the group project in Stages 1 and 2, and those who have not yet chosen the group project may then switch to the group project.

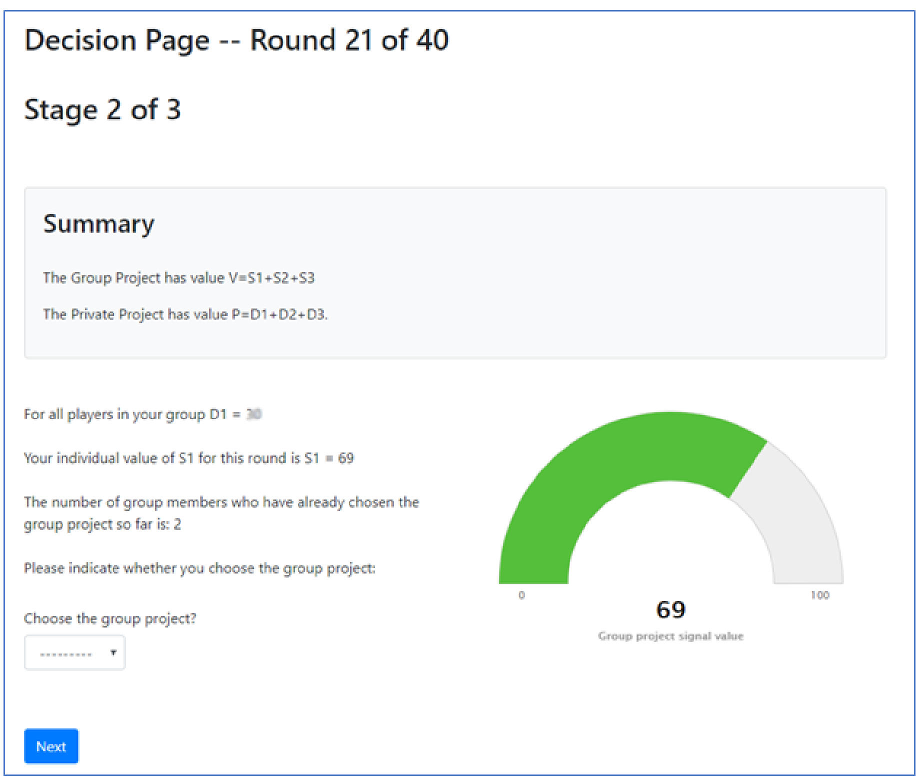

Figure 7: Second Stage to Choose the Group Project

Note: The above *Three Choice Stages* subsection was included for only the Dynamic treatment. In the Static treatment this was replaced with the following paragraph, and Figure 7 was omitted.

You and the others in your group, will make your decisions at the same time. In other words, everyone in your group makes their choice before learning the choices of other group members.

End of the Period

After all members of your group have made their choices, you will learn the values of the group project (V) and the private project (P), and your earnings in experimental dollars for the period. You will also learn how many other members of your group chose the group project, and the other two D2 and D3 draws that determine the private project value P.

As illustrated in Figure 8, your computer will also display at the end of the period a summary of the results from all previous periods in this part of the experiment, in a table you can scroll through if desired.

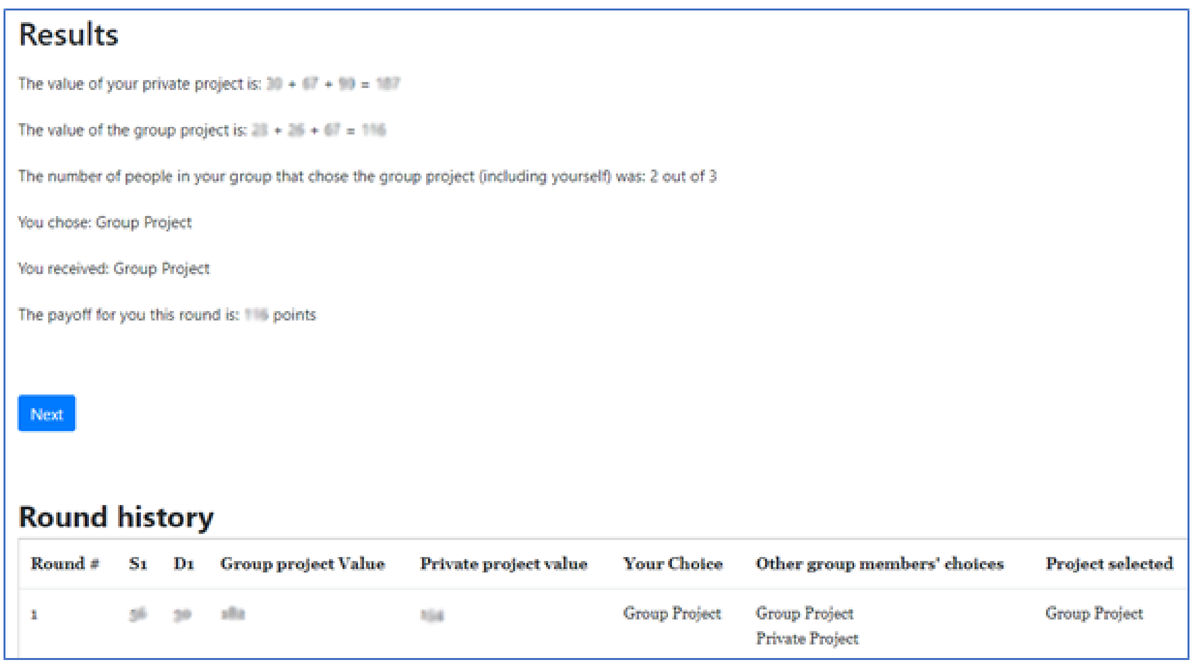

Figure 8: Results Screen

Remember that you will be randomly and anonymously re-matched into new groups of three at the start of each period. Also remember that signals for the group project and the D2 and D3 draws for the private project are randomly and independently drawn for each member of your group.

Of the 60 periods in Part 1, one will be randomly selected for payment. All participants will be paid their earnings converted to US dollars for the randomly selected period, plus a $5.00 show-up payment. You will not find out which period you will be paid for until the end of the experiment, so you should treat each period as something for which you might get paid. You will not be paid for the periods that are not randomly selected for payment.

Summary of Part 1A

In each period:

- The value of the group project is given by V = S1 + S2 + S3. You will observe your own signal, but not the signals of the other members of your group.
- The value of the private project is given by P = 70 + D2 + D3. You will observe the baseline value 70 but not D2 or D3.
- **Static treatment only: You and others in your group make your choice for the group project or private project at the same time, before learning the choices made by any other group members.**
- **Dynamic treatment only: You may choose the group project in one of three stages. After choosing the group project in a period you cannot switch back to choose the private project. But if you do not choose the group project in the first stage, then you may do so in the second stage. If you do not choose the group project in the first or second stage, you may do so in the third stage. At each stage, you will observe how many of your group members chose the group project in a prior stage.**
- If you chose the group project, and at least one other member of your group also chose the group project, then you will earn V (the value of the group project).
- If you do not choose the group project, or you are the only member of your group who chose the group project, then you will earn P (the value of the private project).
- At the start of each period, you will be randomly and anonymously matched into groups of three. At the start of each later period, you will be randomly and anonymously re-matched into new groups of three and you never learn the identities of the other group members in

any period. It is possible, but unlikely, that you may be grouped with the same people in two consecutive periods.

Are there any questions before we begin the experiment?

Distributed separately at the end of the session:

EXPERIMENT INSTRUCTIONS PART TWO

Part 2 will consist of two periods of decisions. You will be paid, in experimental dollars, for the sum of your earnings in both periods. At the end of the experiment we will convert the experimental dollars you earn in part 2 to U.S. dollars at an exchange rate of 50 experimental dollars equals $1.

In each period, the computer will randomly draw an integer from a pre-specified interval. The interval with either be from 0 to 99, or from 20 to 129. Each number in the interval will be equally likely to be chosen. In each period, you will be required to submit a bid to the computer.

If your bid is greater than or equal to the random number, you will receive 100 experimental dollars, plus 1.5 times the random number, minus your bid. If your bid is less than the random number you will receive 100 experimental dollars.

If, for example, you bid 42:

- Suppose the value of the random number is 36. Then your payoff will be 100 + 1.5*36 - 42 = 112.
- Suppose the value of the random number is 20. Then your payoff will be 100 + 1.5*20 - 42 = 88.
- Suppose the value of the random number is 67. Then your payoff will be 100.

Your results from each period will be displayed on the screen.

Are there any questions before we begin part 2?

B Theory

For ease of exposition we present the game with a continuous signal space on the interval $[0, 1]$. In our experimental implementation we use a discrete signal space on the interval $[0, 100]$ to avoid the need to use decimal notation, and we confirm the discrete equilibrium via numerical methods.

B.1 Static treatment

Consider first the static treatment. As discussed in the text, signals and value components are drawn independently from a uniform distribution. Therefore, an agent who receives the private good will earn, in expectation, $\mathbb{E}[V_i] = \mathbb{E}[D_0 + D_{1,i} + D_{2,i}] = D_0 + 1$. Ex-ante, and ignoring any potential selection effects, an agent who receives the public good will earn, in expectation, $\mathbb{E}[P] = \mathbb{E}[s_1 + s_2 + s_3] = s_1 + 1$. As noted in the main text, this comparison suggests the simple but incorrect decision rule of selecting the PG if and only if $s_1 \geq D_0$, as suggested by cursed equilibrium (Eyster and Rabin, 2005), in which agents make the erroneous assumption that other players' decisions are not conditioned on their private information.

A (symmetric) equilibrium consists of two pieces of information: a probability distribution $f_i(s_i|S, D_0)$ that specifies the probability that agent i selects the PG given a private signal s_i in the static treatment (S) with initial private value D_0, and a belief function $\beta_i(s_j|S, D_0)$ which represents the belief of agent i that another player j will select the PG given a private signal s_j. In a Nash equilibrium beliefs must be correct, such that $\beta_i(s_j|S, D_0) = f_j(s_j|S, D_0)$. In a cursed equilibrium beliefs are correct on average, but independent of s_j, such that $\beta_i(s_j|S, D_0) = \int_{s=0}^{1} f_j(s|S, D_0)ds$ for all $s_j \in [0, 100]$.

Given that the expected value of the PG is strictly increasing in s_i, and the value of the private good is constructed to be independent of s_i, an agent who is maximizing their expected payoff, conditional on beliefs, will use a cutoff strategy:

$$f_i(s_i|S, D_0) \begin{cases} = 0 & \text{if } s_i < y_i \\ \in [0, 1] & \text{if } s_i = y_i \\ = 1 & \text{if } s_i > y_i \end{cases}$$

Therefore, we can summarize each agent's behavior by their cutoff point, y_i, where the agent selects the PG when $s_i > y_i$ and selects the private good when $s_i < y_i$.

B.1.1 Nash equilibrium

Writing $\hat{y}_i$ for player i's belief about the value of the cutoff used by others, the expected value of selecting the PG in a Nash equilibrium is given by:

$$\hat{y}_i^2\mathbb{E}[V_i] + (1-\hat{y}_i)^2(s_i + \frac{1+\hat{y}_i}{2} + \frac{1+\hat{y}_i}{2}) + 2(1-\hat{y}_i)(\hat{y}_i)\Big(s_i + \frac{1+\hat{y}_i}{2} + \frac{\hat{y}_i}{2}\Big).$$

The first term is the payoff when no one else selects the PG and the private good is realized, the second term is the payoff in the case where both other players select the PG and the third term is the case where exactly one other player selects the PG.

Therefore, an agent is indifferent between selecting the PG and the private good whenever

$$\mathbb{E}[V_i] = \hat{y}_i^2\mathbb{E}[V_i] + (1-\hat{y}_i)^2(s_i + \frac{1+\hat{y}_i}{2} + \frac{1+\hat{y}_i}{2}) + 2(1-\hat{y}_i)(\hat{y}_i)\Big(s_i + \frac{1+\hat{y}_i}{2} + \frac{\hat{y}_i}{2}\Big) \iff$$
$$(1-\hat{y}_i^2)(D_0+1) = (1-\hat{y}_i)^2(s_i + \frac{1+\hat{y}_i}{2} + \frac{1+\hat{y}_i}{2}) + 2(1-\hat{y}_i)(\hat{y}_i)\Big(s_i + \frac{1+\hat{y}_i}{2} + \frac{\hat{y}_i}{2}\Big).$$

This implies, after some algebra, that the agent is indifferent between the PG and the private good when their signal satisfies the following equality:

$$s_i = D_0 + 1 - \frac{1-\hat{y}_i^3}{1-\hat{y}_i^2}$$

A (symmetric) Bayesian Nash equilibrium cutoff, denoted by y^N, where agents select the PG with a positive probability, is therefore given by:

$$y^N = D_0 + 1 - \frac{1-{y^N}^3}{1-{y^N}^2}$$

B.1.2 Cursed equilibrium

In a Cursed equilibrium each agent (erroneously) believes that others are making decisions that are not dependent on their privately observed signals. We shall denote the belief about this fixed probability by $\bar{y}$. Therefore, in expectation, the PG generates a payoff of:

$$\bar{y}_i^2\mathbb{E}[V_i] + (1-\bar{y}_i^2)(s_i + \frac{1}{2} + \frac{1}{2}),$$

where the first term is the payoff when no one else selects the PG and the private good is realized, and the second term is the payoff in the case where the PG is realized.

Therefore, an agent is indifferent between selecting the PG and the private good whenever

$$\begin{aligned}
\mathbb{E}[V_i] &= \bar{y}_i^2\mathbb{E}[V_i] + (1-\bar{y}_i^2)(s_i + \frac{1}{2} + \frac{1}{2}) \iff \\
(1-\bar{y}_i^2)\mathbb{E}[V_i] &= (1-\bar{y}_i^2)(s_i + \frac{1}{2} + \frac{1}{2}) \iff \\
\frac{(1-\bar{y}_i^2)(D_0+1) - (1-\bar{y}_i^2)}{1-\bar{y}_i^2} &= s_i \iff \\
D_0 &= s_i.
\end{aligned}$$

Therefore, the Cursed Equilibrium cutoff, y^C, is given by $y^C = D_0$. Table 6 summarizes these cutoffs for the three values of D_0 employed in the experiment.

	Nash Equilibrium Cutoffs	Cursed Equilibrium Cutoffs
$D_0 = 0$	$y^N = 0$	$y^C = 0$
$D_0 = 0.3$	$y^N = 0.250$	$y^C = 0.3$
$D_0 = 0.7$	$y^N = 0.521$	$y^C = 0.7$

Table 6: Equilibrium cutoffs in the Static treatment.

B.2 Dynamic Treatment

The dynamic treatment employs three simultaneous stages of decision making. In each stage, agents decide whether to select the PG or not. A selection of the PG is binding, but a choice not to select the PG is a deferral of decision-making to subsequent stages. There exists, therefore, an equilibrium where all agents defer decision making in the first two stages and then play the static equilibrium in the final stage. There also exists an equilibrium where all agents defer decision making in the first stage and only choose the PG with a positive probability in the final two stages. Neither of these equilibria are particularly intuitive, as they lead to lower expected payoffs than in an equilibrium where agents with strong signals select the PG in the first stage.[17]

[17] In addition, there is no evidence in the data that these delayed equilibrium are being used. Of the 144 subjects, 143 selected the PG at least once in the first stage, and 66 percent of all PG choices in the dynamic treatment

The dynamic treatment has four distinct observable decision making states. First, the decision in the first stage. Second, the decision in the second stage conditional on observing both other agents selecting the PG in the first stage. Third, the decision in the second stage conditional on observing exactly one other agent selecting the PG in the first stage. Fourth, the decision in the third stage conditional on observing one agent select the PG in the first stage, and the other agent select the PG in the second stage. In equilibrium, if no agent actively selects the PG in a particular stage, then no agent will select the PG in any later stage. The reasoning is straightforward: if an agent does not wish to select the PG now, then they will not wish to select the PG in any later stage unless they observe a positive signal about the value of the PG (i.e., unless they observe someone else selecting the PG before the next stage begins).

Denote the four cutoff values by y_0, y_1, y_2, $y_{1,1}$, denoting the first stage decision, the case where one other person selected the PG in the first stage, the case where both others selected the PG in the first stage, and the third stage case where one other selected the PG in each of the first two stages, respectively.

B.2.1 Subgame Perfect Nash equilibrium

As discussed above, we focus our attention to the SPNE in which any agent with a high enough signal selects the PG in the first stage.

Using standard notation, we shall refer to the equilibrium cutoff values by y_0^*, y_1^*, y_2^* and $y_{1,1}^*$, respectively. We proceed via backwards induction. There are two situations where an agent essentially faces an individual decision problem: the case where both other players have chosen the PG in the first period (so y_2^* is relevant), and the case where one agent chooses the PG in the first stage and one chooses the PG in the second (so $y_{1,1}^*$ is relevant).

We begin with the case where an agent has declined to select the PG in the first two stages and observed one player select the PG in the first stage and the other select the PG in the second stage. In this case, it must be that $y_0^* \geq y_1^*$; otherwise it is not possible to observe one agent choosing the PG in each of the first two stages. In the third stage, the agent's expected value for the PG (which she will receive with certainty if she selects the PG) is given by:

$$s_i + \frac{1 + y_0^*}{2} + \frac{y_0^* + y_1^*}{2}.$$

occurred in the first stage.

As always, the expected value of the private good is $D_0 + 1$. Therefore, for this agent's point of indifference between the PG and the private good is:

$$s_i + \frac{1 + y_0^*}{2} + \frac{y_0^* + y_1^*}{2} = D_0 + 1 \iff$$
$$s_i = D_0 + \frac{1}{2} - y_0^* - \frac{y_1^*}{2}$$

whenever $0 \leq D_0 + \frac{1}{2} - y_0^* - \frac{y_1^*}{2}$. If $0 > D_0 + \frac{1}{2} - y_0^* - \frac{y_1^*}{2}$ then the agent *always* prefers the PG. Therefore,

$$y_{1,1}^* = \max\{0, D_0 + \frac{1}{2} - y_0^* - \frac{y_1^*}{2}\}. \tag{1}$$

Next, consider the case where the agent in the second stage has observed both other players select the PG in the first stage. In this case, the expected value of the PG is given by

$$s_i + \frac{1 + y_0^*}{2} + \frac{1 + y_0^*}{2}.$$

In this case, the indifference point between the PG and the private good occurs when $s_i = D_0 - y_0^*$ whenever $D_0 \geq y_0^*$. Therefore

$$y_2^* = \max\{0, D_0 - y_0^*\}. \tag{2}$$

Note that $y_{1,1}^* \geq y_2^*$. This is intuitive, as observing both players choosing the PG immediately is a stronger positive signal about its value, so a rational agent is willing to select the PG at a lower s_i.

Consider next the case where one agent has chosen the PG in the first stage, to determine y_1^*. The remaining two agents then play out a two player continuation game, where each agent believes that, in expectation, the signal of the player who already chose the PG is $\frac{1+y_0^*}{2}$. Two cutoff points are relevant in this continuation game: y_2 and $y_{1,1}$.[18]

There are two cases to consider. In the first case, suppose that $y_{1,1}^* \geq s_i$ so that the agent will not select the PG in the third stage (if they have not already done so in the second stage) even if

[18] We maintain our assumption of focusing only on equilibrium without delay, so that if no one selects the PG in the second stage then they will not select it in the third stage either.

they observe an entry in the second stage. Therefore, in this case, an agent who does not select the PG in the second stage will receive, in expectation, $D_0 + 1$. The expected value of selecting the PG in the second stage is given by $s_i + \frac{1+y_0^*}{2} + \frac{y_0^*}{2}$. In any equilibrium, the agent must be indifferent between these two options when $s_i = y_1^*$, which implies that $y_1^* = D_0 + \frac{1}{2} - y_0^*$. But, this implies that $y_1^* > y_{1,1}^*$, which violates our case condition. Therefore, there is no equilibrium where $y_1^* \leq y_{1,1}^*$.

In the second case we have $s_i > y_{1,1}^*$, such that an agent who defers a decision to the third stage will select the PG if and only if the remaining player selects the PG in the second stage. The critical event is the case where the other agent does not select the PG in the second stage: if the other player does select the PG, then the agent will select the PG in the third stage (rendering the second stage decision moot). In this critical event, the payoff from selecting the PG in the second stage is $s_i + \frac{1+y_0^*}{2} + \frac{y_1^*}{2}$, and the payoff from not selecting the PG is $D_0 + 1$. Therefore, the equilibrium indifference point must satisfy

$$y_1^* + \frac{1+y_0^*}{2} + \frac{y_1^*}{2} = D_0 + 1 \iff$$

$$y_1^* = \frac{2}{3}[D_0 + 1 - \frac{1+y_0^*}{2}] \tag{3}$$

and substituting into the expression for $y_{1,1}^*$ yields

$$y_{1,1}^* = \max\{0, y_1^* - \frac{y_0^*}{2}\} \tag{4}$$

Equilibrium behavior in the continuation game that begins in the second stage after one player has selected the PG in the first stage is summarized by Equation 2, Equation 3, and Equation 4.

Next, we briefly consider the continuation game where no agent selected the PG in the first stage. The natural equilibrium in this continuation game is for no agent to select the PG in the second or third stage either; after observing no one selecting the PG in the first stage, given the use of cutoff strategies, the expected value of the PG must weakly decrease. Therefore, the expected payoff in this continuation game is $D_0 + 1$ for all agents.

Finally, consider behavior in the first stage. There are two key cases. In the first, $y_1^* > s_i$, and in the second, $s_i \geq y_1^*$.

In the first case the agent i observes a signal such that $s_i < y_1^*$, and suppose that an equilibrium exists in which agent i selects the PG in the first stage. There are three sub-cases. In the first sub-case, both other agents select the PG in the first stage and agent i is indifferent between selecting or not selecting the PG in the first stage (if she wishes to receive the PG, she can simply select the PG in the second stage). In the second sub-case, exactly one other agent selects the PG in the first stage. In this case, agent i now prefers to not receive the PG (because $y_1^* > s_i$). In the third sub-case, both other agents decline the PG in the first stage. Because $s_i < y_1^*$, and agent i selected the PG, it must be that $y_0^* < y_1^*$, so that the other agents will not select the PG in the second stage and, therefore, agent i is indifferent between having selected the PG or not in the first stage (as she receives the private good in either case). Pulling the three sub-cases together, agent i is either indifferent or strictly prefers *not* to have selected the PG in the first stage. Therefore, agent i cannot select the PG in any equilibrium when $s_i < y_1^*$. There does not exist any equilibrium with $y_0^* < y_1^*$.

Now consider the second case, where $s_i \geq y_1^*$. As before, the critical decision occurs when no other agents select the PG in the first stage: if another agent selects the PG immediately, then an agent can still secure the PG by selecting it in the next stage.

If agent i does not select the PG in the first stage, conditional on no other agents selecting the PG in the first stage, then the PG will not be provisioned in any non-"delay" equilibrium. Therefore, the expected payoff for all agents is $D_0 + 1$. If agent i does select the PG in the first stage, then the outcome is determined by the behavior of the other agents in the second stage. With probability $\frac{y_1^{*2}}{y_0^{*2}}$ the PG is not provisioned, and the expected payoff from selecting the PG in the first stage is $\frac{y_1^{*2}}{y_0^{*2}}(D_0+1) + \frac{2y_1^*(y_0^*-y_1^*)}{y_0^{*2}}(s_i + \frac{y_1^*}{2} + \frac{y_0^*+y_1^*}{2}) + \frac{(y_0^*-y_1^*)^2}{y_0^{*2}}(s_i + y_0^* + y_1^*)$.

The agent is indifferent when the private signal about the value of the private good satisfies:

$$(y_0^{*2} - y_1^{*2})(D_0+1) = 2y_1^*(y_0^* - y_1^*)(s_i + \frac{y_1^*}{2} + \frac{y_0^* + y_1^*}{2}) + (y_0^* - y_1^*)^2(s_i + y_0^* + y_1^*),$$

which implies that, in equilibrium, the following implicit condition for y_0^* must hold:

$$(y_0^{*2} - y_1^{*2})(D_0+1) = 2y_1^*(y_0^* - y_1^*)(y_0^* + \frac{y_1^*}{2} + \frac{y_0^* + y_1^*}{2}) + (y_0^* - y_1^*)^2(y_0^* + y_0^* + y_1^*). \tag{5}$$

Solving Equation 5 and Equation 3 simultaneously (and selecting the solution that ensures $1 > y_0^* > y_1^* \geq 0$) generates the equilibrium. Table 7 summarizes the four key cutoffs for the D_0 values used in the experiment.

	Nash Equilibrium Cutoffs				Cursed Equ.
	y_0^*	y_1^*	$y_{1,1}^*$	y_2^*	y^C
$D_0 = 0$	0.477	0.175	0	0	0
$D_0 = 0.3$	0.589	0.337	0.043	0	0.3
$D_0 = 0.7$	0.730	0.557	0.192	0	0.7

Table 7: Equilibrium cutoffs in the dynamic treatment.

References

Bagnoli, M. and Lipman, B. L. (1989). Private provision of public goods: Fully implementing the core through private contributions. *The Review of Financial Studies*, 56(4):583–601. → pages 7

Bazerman, M. H. and Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resoution*, 27(4):618–634. → pages 1

Capen, E., Clapp, R., and Campbell, W. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology*, 23:641–653. → pages 1

Charness, G. and Levin, D. (2009). The origin of the winner's curse: A laboratory study. *American Economic Journal: Microeconomics*, 1(1):207–236. → pages 1

Chen, D. L., Schonger, M., and Wickens, C. (2016). otree – an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97. → pages 11

Cox, C. (2015). Cuesed beliefs with common-value public goods. *Journal of Public Economics*, 121:52–65. → pages 2, 5

Cox, C. and Stoddard, B. (2021). Common-value public goods and informational social dilemmas. *American Economic Journal: Microeconomics*, page forthcoming. → pages 5

Dickinson, D. (1998). The voluntary contributions mechanism with uncertain grou payoffs. *Journal of Economic Behavior and Organization*, 35:517–533. → pages 5

Esponda, I. and Vespa, E. (2014). Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 6(4):180–202. → pages 1, 2, 3

Esponda, I. and Vespa, E. (2019). Contingent preferences and the sure-thing principle: Revisiting classic anomalies in the laboratory. → pages 3, 4

Eyster, E. (2019). Chapter 3 - errors in stratgeic reasoning. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:187–259. → pages 1

Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672. → pages 1, 7, 34

Gangadharan, L. and Nemes, V. (2009). Experimental analysis if risk and uncertainty in provisioning private and public goods. *Economic Inquiry*, 47:146–164. → pages 5

Gillen, B., Snowberg, E., and Yariv, L. (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy*, 127(4):1826–1863. → pages 13, 21

Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125. → pages 11

Kagel, J. H. and Levin, D. (1986). The winner's curse and public information in common value auctions. *American Economic Review*, 76:894–920. → pages 1

Kagel, J. H. and Levin, D. (2002). *Common Value Auctions and the Winner's Curse*. Princeton University Press, Princeton, NJ. → pages 1

Levati, M. V., Marone, A., and Fiore, A. (2009). Voluntary contributions with imperfect information: An experimental study. *Public Choice*, 138:199–216. → pages 5

Li, S. (2017). Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–3287. → pages 3

Martinez-Marguina, A., Niederle, M., and Vespa, E. (2019). Failures in contingent reasoning: The role of uncertainty. *American Economic Review*, 109(10):3437–3474. → pages 3

Ngangoue, M. K. and Weizsacker, G. (2021). Learning from unrealized versus realized prices. *American Economic Journal: Microeconomics*, page forthcoming. → pages 1, 3

Oprea, R. (2020). What makes a rule complex? *American Economic Review*, 110(12):3913–3951. → pages 3

Stoddard, B. (2017). Risk in payoff-equivalent appropriation and provision games. *Journal of Behavioral and Experimental Economics*, 69:78–82. → pages 5