UMass Lowell Economics Working Paper No. 3021

Norm Enforcement with Incomplete Information

Lawrence R. De Geest
Suffolk University

David C. Kingsley
University of Massachusetts Lowell

November 2019

Department of Economics
University of Massachusetts Lowell
http://www.uml.edu/economics
Twitter: @uml_econ Phone: 978-934-2780

Norm Enforcement with Incomplete Information

Lawrence R. De Geest*[1] and David C. Kingsley[2]

[1]Department of Economics, Suffolk University
[2]Department of Economics, University of Massachusetts Lowell

Abstract

We study the emergence of norms and their enforcement in a public goods game with private information about endowments. Subjects were randomly assigned a *Low* or *High* endowment and across treatments endowments were either *Observed* or *Unobserved*. We estimate contribution norms and then estimate the expected costs of noncompliance. We find that incomplete information does not affect norms, but rather their enforcement. In both *Observed* and *Unobserved* we see a "contribute-your-endowment" norm emerge. Enforcement in *Observed* is close to theoretical predictions. However, enforcement in *Unobserved* depended on how well subjects could map contributions to endowments in a given round. When at least one *High* type pooled with *Low* types (by contributing less than or equal to the *Low* endowment), punishment was used to protect *Low* rather than attack *High*: contributions equal to the *Low* endowment were not punished (in case they came from a cooperative *Low* type) while contributions of zero were punished as if they were from a *High* type. This kept cooperation from unraveling, but it also enabled *High* types to hide behind small endowments. Our results dovetail with results from bargaining games and suggest that in settings with incomplete information, norms emerge to attenuate rather than eliminate non-cooperative behavior.

Keywords: Norms; Income Inequality; Incomplete Information; Cooperation; Punishment; Public Goods
JEL Codes: C92, H41, D82

*Corresponding author. Thanks for helpful comments from Brock Stoddard, Alex Smith, Darlene Chisholm, Mackenzie Alston and participants from the 2019 New England Experimental Economics Workshop and the 2019 Asia Pacific Meeting of the Economics Science Association. Funding from University of Massachusetts Lowell is gratefully acknowledged. Declarations of interest: none.

1 Introduction

Splitting a surplus fifty-fifty is an established norm (Krupka and Weber, 2013; Andreoni and Bernheim, 2009), allowing proposers in ultimatum games with private information to "hide behind small cakes" (Güth et al., 1996; Mitzkewitz and Nagel, 1993). The proposer with a large cake offers a fifty-fifty split of a small cake, and the responder – unaware of the true cake size and unwilling to accidentally punish a fair offer from someone who simply has less to share – accepts.

We show that a similar pattern of behavior plays out in a more complex setting: a linear public goods game with peer punishment and private information. We create private information through heterogeneous endowments. Subjects were split into groups of four. Two group members received a high endowment of 30 (*High* types) and the other two a low endowment of 10 (*Low* types). In our control (*Observed*), subjects had complete information and could observe both the contributions and the endowments of each group member. In our treatment (*Unobserved*), subjects had incomplete information and could only observe contributions although they knew the distribution of endowments across group members.

We infer norms using a modified version of the Carpenter and Matthews (2009) contribution norms model. Contributions below some amount – the norm – are more likely to be punished (extensive margin) and punished more severely (intensive margin). The model searches over the set of feasible contributions for each endowment to identify the norms for each margin. We then use the estimated norms and observed punishment between subjects to estimate the expected costs of noncompliance.[1] Our approach not only tells us *what* norms emerged, but also *how* they emerged.

The results from *Observed* replicate the main findings in the literature. The contribution norm was to contribute your entire endowment, just like in games with homogeneous endowments (Carpenter and Matthews, 2009; Nicklisch and Wolff, 2011) and with heterogeneous endowments with complete information (Reuben and Riedl, 2013). Both *Low* and *High* types enforced the norm. Moreover, the expected costs for noncompliance are close to theoretical predictions.

In *Unobserved* we also see a "contribute-your-endowment" norm emerge. However, enforcement of the norm varied with information. This is because period-by-period information in *Unobserved* was endogenous. *High* types could either "conceal" (contribute 10 or less) or "reveal" (contribute more than 10) – but only for that period, because IDs in the punishment stage were randomized. We account for this endogeneity of information in our analysis.

When at least one *High* type concealed, the enforcement rule was to protect *Low* types rather

[1] There are several ways to study injunctive norms (norms that dictate what action(s) people should take). Krupka and Weber (2013) measure *beliefs* about norms using a coordination game in which subjects independently rate the appropriateness of all possible splits in a dictator game and are paid if their rating matches the modal choice. The authors then fit those norms (elicited among one pool of subjects) to data from dictator games (played by another pool of subjects). Kimbrough and Vostroknutov (2016) and Kimbrough and Vostroknutov (2018) introduce rule-following tasks to measure *propensities* for norm compliance and show that norm sensitivity explains individual variation of pro-sociality across experimental settings. We chose to infer norms from punishments for two reasons: because punishments directly reveal what behavior subjects would tolerate, and because we can use the inferred norms and punishments between subjects to estimate the expected costs of noncompliance with norms. That said, our view is that different approaches to measuring norms are probably complements rather than substitutes, and in our discussion we explore how future research can combine these approaches.

than attack *High* types: contributions of ten were not punished, in case they came from cooperative *Low* types, while contributions of zero were punished as if they came from *High* types.[2] Moreover, this enforcement was mostly carried out by *Low* types, while *High* types disengaged. When both *High* types revealed and groups had complete information for that period, there was less punishment on *Low* and more punishment on *High*. But even in these periods with complete information, the punishment on *Low* and *High* was higher than in *Observed*.

There is no doubt that private information imposed social costs. *Low* types were always hit with more punishment in *Unobserved*, and so were *High* types when they revealed. The damage mostly fell to *Low* types: they earned significantly lower payoffs in *Unobserved* than *Observed*. Nevertheless, the norms and enforcement that emerged in *Unobserved* sustained cooperation well above freeriding. At the end of our results we sketch a simple evolutionary model to explore the idea that it is better to allow some bad behavior than try to completely eliminate it.

We make several contributions. For starters, we are the first to estimate norms in a public goods game with private information. We advance previous work on estimating norms from punishment data (Reuben and Riedl, 2013; Nicklisch and Wolff, 2011; Carpenter and Matthews, 2009) by introducing a framework for estimating the expected cost of noncompliance.

Second, we clarify the effects of incomplete information in public goods experiments with peer punishment. Bornstein and Weisel (2010) show that peer punishment is less effective when subjects have private information about their endowments, but their experiment randomly assigns endowments each period, and neither norms nor the expected cost of noncompliance are estimated. We show that information does not affect norms, it affects how norms are enforced. A "contribute-your-endowment" norm emerges in *Unobserved*, but its enforcement allowed *High* types to hide behind small endowments. *High* types received less punishment over the contribution range $[0, 10]$ in *Unobserved* than *Observed*, while *Low* types suffered more punishment. This helps explain why low-endowment subjects prefer enforcement from a central authority when high-endowment subjects have private information (De Geest and Kingsley, 2019).

Third, our results contribute to the literature on punishment in public goods games with information asymmetries. Several studies look at punishment under imperfect information, such as when contributions are observed with noise (Nicklisch et al., 2016; Ambrus and Greiner, 2012; Grechenig et al., 2010). The main difference between incomplete and imperfect information is that subjects cannot exploit noise the way they can exploit private information: a *High* type cannot hide behind a small endowment if there is a chance their "hiding contribution" is flipped to a "revealing contribution". While punishment leads to an unraveling of cooperation in games with imperfect information, we show that punishment produces stable cooperation under incomplete information, albeit less than under complete information.

Looking at the big picture, our results dovetail with results from bargaining games and show a clear pattern of behavior in both complex and simple strategic settings with private information.

[2] Similarly, Ali and Miller (2016) show that a mix of punishment and forgiveness sustains cooperation when groups have incomplete information about individual actions.

Agents *with* private information hide behind small cakes (or small endowments), while agents *without* private information – wary of mistakenly punishing fair offers or contributions – enforce only a minimum standard of cooperation.[3] It seems that the role of norms in settings with private information is to attenuate rather than eliminate bad behavior. At the same time, norms and compliance are sensitive to context (Kimbrough and Vostroknutov, 2016; Krupka and Weber, 2013), so it would be interesting to extend our design to explore how and why norms emerge in other settings with private information. We consider a few ideas in our discussion.

2 Experiment design and methods

We ran a linear public good experiment with endowment heterogeneity and peer punishment. Each subject was randomly allocated a fixed endowment of experimental dollars (EDs) which could be allocated between a private account and a group account. Each group was composed of two *High* endowment members who received 30 EDs and two *Low* endowment members who received 10 EDs. Randomly assigned endowments were maintained for the entire experiment (i.e. once a *High* type always a *High* type). The distribution of group endowments was identical (two *Low* and two *High*) and known (subjects knew they were in a group of two *Low* and two *High*). Groups remained fixed throughout the experiment. The experiment lasted for 50 periods to provide ample time for contribution norms to emerge and for groups to realize the benefits of punishment (Gächter et al., 2008).

Payoffs to subject i were

$$\pi_i = \max\left[0, (e_i - x_i) + \alpha \sum_{j=1}^{n} x_j - r \sum_{j \neq i}^{n-1} P_{ji}\right] - c \sum_{j \neq i}^{n-1} P_{ij} \tag{1}$$

where x_i is the subject's contribution to the group account, e_i is the subject's endowment, $\alpha = 0.4$ is the marginal per capita return (MPCR) from the public good, and $\sum_{j=1}^{n} x_j$ represents the sum of contributions to the group account from all group members. With n players, $\frac{1}{n} < \alpha < 1$, and a known last period, there is a unique, symmetric Nash equilibrium where everybody freerides and contributes nothing to the public good. Similarly, there is a social optimum where subjects contribute their entire endowment to the public good.

Subjects were informed of their endowments at the start of the experiment. At the beginning of each period, each subject chose a contribution to the group account. After all contribution decisions were made, each subject was given the opportunity to punish their group members. In Equation 1 P_{ij} represents the number of reduction points that i imposes on other group members j at a cost of $c = 1$, and P_{ji} represents the number of reduction points that other group members j impose upon i at a cost of $r = 3$. In order to avoid excessive losses and ensure *High* types did not have more power in enforcement than *Low* types, subjects of all endowments were allowed to impose up to 10

[3]Similarly, Nicklisch et al. (2016) study punishment with imperfect information and show that in a very noisy setting, a central authority will abstain from punishing to avoid misguided punishment of cooperative group members.

reduction points per period and losses on any given period were bounded at zero unless the subject imposed punishment (Gächter et al., 2008; Reuben and Riedl, 2013).[4] The costs associated with imposing *reduction points* were referred to as *administrative costs*, while the costs associated with receiving *reduction points* were referred to as *reduction costs*

In the punishment stage subjects were shown: the aggregate contribution to the group account; the individual contributions of their group members by random ID; their individual period earnings; their total earnings (equal to the sum of their individual period earnings); and a history of outcomes in previous periods. The random ID and the order of presentation of the contributions of one's group members was randomized each period to avoid reputation effects. This allowed *High* types to switch between "revealing" (contribute more than 10) and "concealing" (contributing less than 10) and thus made period-by-period information endogenous. We account for this endogeneity in our analysis. Our treatments vary whether subjects could observe endowments alongside individual contributions of group members. In the control, *Observed*, similar to Reuben and Riedl (2013), subjects could link individual contributions to individual endowments. In the treatment, *Unobserved*, subjects could only view contributions. However, a single contribution could be linked to an endowment if that subject contributed more than 10 (thus revealing they are a *High* type). In addition, all contributions could be linked to endowments if both *High* types contribute more than 10, in which case the group has complete information (but only for that period).

2.1 Expected costs of noncompliance

In our analysis we estimate contribution norms in each treatment. We then calculate the expected costs of noncompliance to these norms for *Low* and *High* types controlling for their respective norms. Comparing estimated noncompliance costs to theoretical noncompliance costs shows us whether subjects applied deterrent penalties.

With complete information we expect a "contribute-your-endomwent" norm (the social optimum) to emerge (Reuben and Riedl, 2013; Carpenter and Matthews, 2009). In our design when subjects contribute this norm they earn 32 EDS. Assuming subjects tradeoff the benefits and costs of complying with norms, noncompliance is eliminated when the expected costs are greater than the expected benefits. Figure 1 shows the deterrent expected costs of noncompliance to a *Low* or a *High* type. For each type, the marginal incentive to noncompliance is equivalent: they can increase their individual payoffs $1 - \alpha = 0.6$ EDs for each ED withheld from the public good. However, the total costs of noncompliance vary with endowments. A *High* type earns 50 EDs by freeriding (contributing zero), so the deterrent penalty is 18 EDs; a *Low* type earns 38 EDs by freeriding, so the deterrent penalty is 6 EDs.

[4] We removed power asymmetries in enforcement to focus on the effect of incomplete information in our design. However, power asymmetries are often a consequence of inequality. We explore this idea in our discussion as a topic for future research.

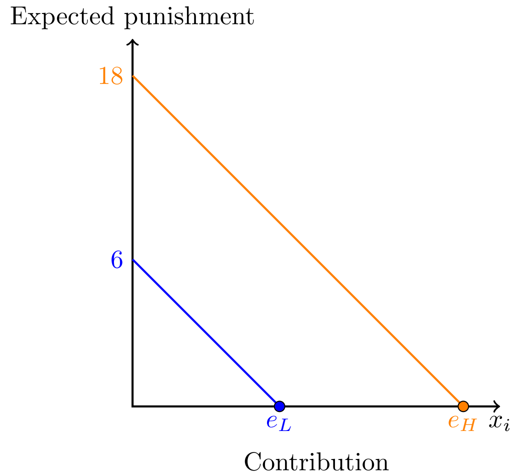

Figure 1: Expected costs of noncompliance with the "contribute-your-endowment" norm (e_L for *Low* and e_H for *High*). The expected cost of freeriding (contribute zero) is 6 EDs for *Low* (38 EDs - 32 EDs) and 18 EDs for *High* (50 EDs - 32 EDs).

2.2 Implementation

We ran the experiment in November and December 2018 at the Cleve E. Willis experimental lab at the University of Massachusetts Amherst. We recruited subjects from the undergraduate population using ORSEE (Greiner, 2015) and implemented the experiment in z-Tree (Fischbacher, 2007). At the beginning of the experiment we passed out and read the instructions. Then we required each participant to correctly answer a set of comprehension questions before the experiment would continue.[5] Across 4 sessions, 9 and 10 groups participated in *Observed* and *Unobserved* respectively. Each session lasted about 60 minutes. Subjects earned on average $17.26, including a $7 show-up fee.

3 Results

We begin by summarizing average contributions, punishments, and earnings in Section 3.1.[6] We find that contributions by *High* went up significantly in *Observed*, and as a result, so too did earnings of *Low*. However, a closer look reveals that contributions by *High* were bi-modal in both treatments, peaking at 10 in *Unobserved* and 30 in *Observed*.

In Section 3.2 we back out the contribution norms in each treatment for *Low* and *High* and then calculate the expected costs of noncompliance. In *Observed* we find a "contribute-your-endowment" norm (similar to Reuben and Riedl, 2013; Nicklisch and Wolff, 2011; Carpenter and Matthews, 2009), and the costs of noncompliance are close to theoretical predictions. But in *Unobserved* the norms and noncompliance costs that emerge suggest subjects tried to strike a balance between ensuring *High* did not freeride and *Low* were not punished for cooperating. We motivate this result with a simple evolutionary model.

[5] Our experiment instructions are in Section D of the appendix.

[6] The code to replicate our analysis can be found at `https://github.com/lrdegeest/NormEnforcement`.

3.1 Contributions, Punishments, and Earnings

Table 1 displays the average group contributions, punishment (sent and received), and earnings (EDs) across treatments, overall and by endowment type. The effect of information on average contributions is immediately clear, with the overall effect being driven by the behavior of *High* types. Contributions are significantly higher overall ($z = 1.96$, $p = 0.05$) and among *High* types ($z = 2.287$, $p = 0.022$) but statistically equivalent among *Low* types ($z = 1.43$, $p = 0.253$).[7]

Table 1: Average Contributions, Punishment and Earnings across Treatments

	Contributions		**Earnings**		**Punishment Sent**		**Punishment Received**	
	Observed	*Unobserved*	*Observed*	*Unobserved*	*Observed*	*Unobserved*	*Observed*	*Unobserved*
Pooled	14.81	9.67	35.87	33.52	0.897	0.788	2.69	2.36
	(5.31)	(5.64)	(5.97)	(4.06)	(0.915)	(0.144)	(2.75)	(1.36)
High	21.40	12.64	38.37	40.41	1.14	0.598	3.32	2.35
	(8.40)	(8.82)	(4.31)	(1.51)	(1.41)	(0.489)	(2.94)	(1.25)
Low	8.21	6.71	33.36	26.22	0.656	0.977	2.06	2.37
	(2.43)	(2.94)	(8.48)	(7.48)	(0.455)	(0.763)	(2.70)	(1.85)

Figure 2 shows average contributions over time (on the left) and the distributions of contributions over time (on the right). The time series of average contributions suggests that contributions were fairly stable after an initial learning phase in the first ten or so periods. So, in displaying the distributions of contributions over time we broke them up by the first ten periods and periods 39-49 (we exclude the last period because of endgame effects). Contributions by *Low* piled up at around 10 (their endowment) in *Observed*, while they were fairly spread out in *Unobserved*. By contrast, the distribution of contributions by *High* types had two peaks, one at 10 and the other at 30. Considering only the later periods, there is a clear variation in the contributions of *High* types across treatments. In *Observed*, most contributions migrated over time towards 30. In *Unobserved*, the density in the middle flattened out and contributions were pushed towards 10 or 30. However, overall contributions by *High* in *Unobserved* were not significantly different from 10 ($z = 0.459$, $p = 0.646$).[8]

[7] All Wilcoxon Ranksum tests are conducted at the group level including 9 and 10 observations in *Observed* and *Unobserved* respectively.

[8] Wilcoxon signed-rank test with 10 group-level observations.

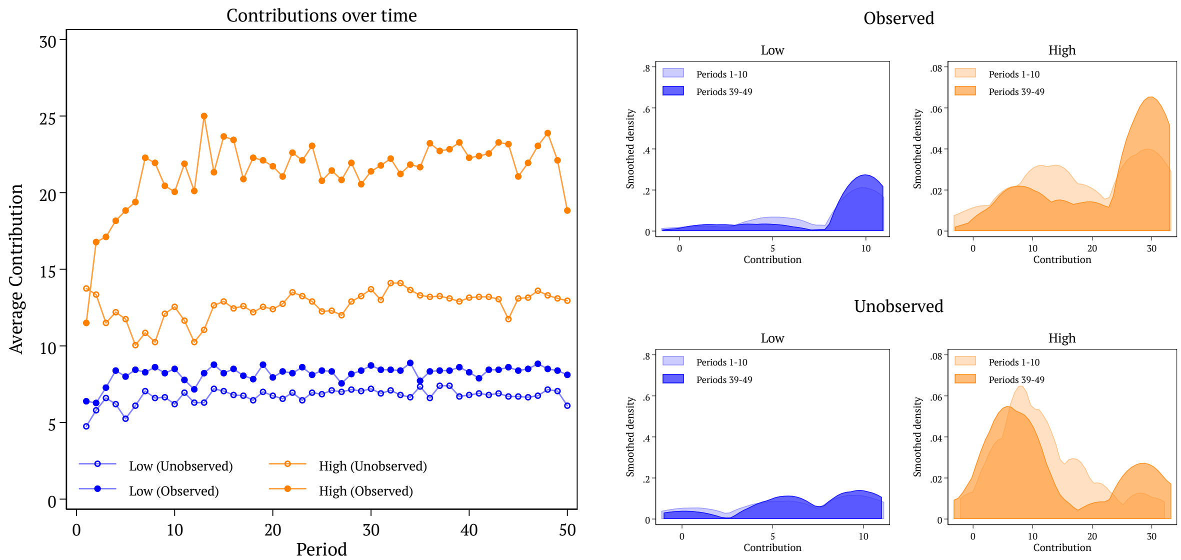

Figure 2: Summary of contributions. The lefthand plot shows average contributions over time broken down by endowment and treatment. The righthand plots show the distributions of contributions. Distributions are broken up between the periods 1-10 and periods 39-49 to exclude end-game effects in the last period.

Turning to punishment sent and received, there is no significant difference in punishment sent across treatments overall ($z = 0.653$, $p = 0.514$), among *High* ($z = 0.00$, $p = 1.00$), or among *Low* ($z = 1.061$, $p = 0.289$). Similarly, there is no significant difference in punishment received across treatments overall ($z = 0.653$, $p = 0.514$) or among *High* ($z = 0.327$, $p = 0.744$). There is a small and marginally significant difference in punishment received among *Low* ($z = 1.677$, $p = 0.094$) suggesting they receive more punishment in *Unobserved* than *Observed*. However, these results combine the extensive and intensive margins of punishment and do not consider deviations from norms. In Section 3.3 we show that treatment differences appear when we separately estimate the margins and then calculate the expected costs of noncompliance to contribution norms.

Finally, there is no significant difference in average group earnings across *Observed* and *Unobserved* ($z = 1.31$, $p = 0.19$). However, looking across endowments reveals that *Low* types in *Observed* earn significantly more than their *Unobserved* counterparts ($z = 1.96$, $p = 0.05$) while *High* types earn a statistically equivalent amount across treatments ($z = 0.653$, $p = 0.51$).

Our results so far show that full transparency pushes contributions up to endowments. While freeriding is not completely eliminated, particularly among *High* types, the proportion of subjects contributing their full endowment are the majority in *Observed*. In *Unobserved* contributions from *Low* also trend towards their endowment. While most *High* types exploit their private information, it is interesting that we see a peak emerge at 30. In the next section we explore how punishment may have shaped this dynamic.

3.2 Contribution norms

Contribution norms are simply the level of contributions subjects expected from each other. Subjects who did not comply with the norm were punished.

To estimate contribution norms we use a modified version of the model introduced by Carpenter and Matthews (2009).[9] In each period, each group member i first decides whether or not to punish each other group member j (the extensive margin), and if so, how much to punish (the intensive margin).

Figure 3 illustrates the basic idea of the model on the extensive margin of punishment. Starting with a simple example, Panel 3a shows a linear probability model in which the probability of punishment depends only on a target's contribution x and how it compares to the contribution norm γ. In the linear model, γ is a kink or discontinuity, with β_1 the slope before the kink and $\beta_1 + \beta_2$ the slope after the kink. So, starting from γ, a marginal *decrease* in the target's contribution leads to a $|\beta_1|$ *increase* in the probability of punishment, while a marginal *increase* in the target's contribution leads to a $\beta_1 + \beta_2$ *decrease* in the probability of punishment, implying $|\beta_1| > |\beta_1 + \beta_2|$.

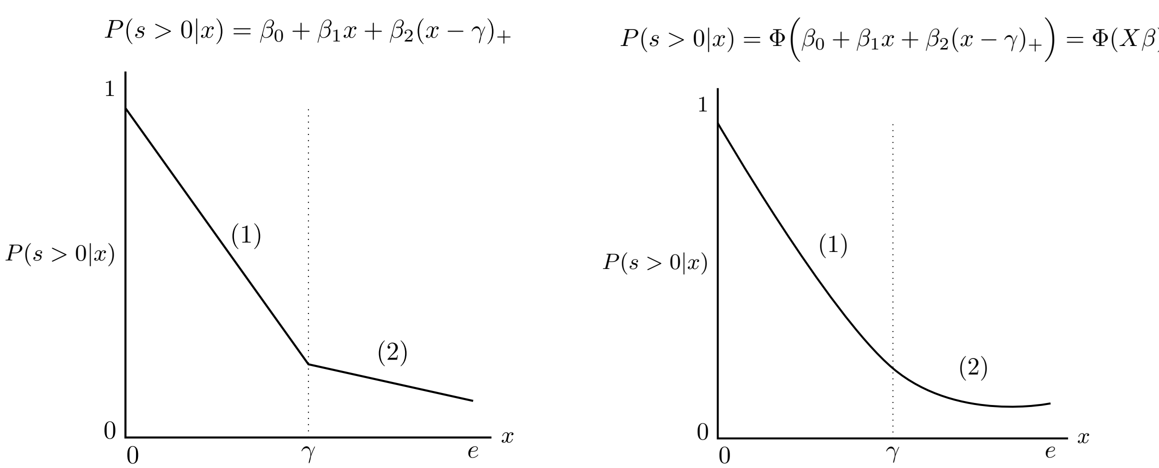

(a) Linear norms model. Marginal punishment is (1) β_1 when $x \leq \gamma$ and (2) $\beta_1 + \beta_2$ when $x > \gamma$.

(b) Nonlinear norms model. Marginal punishment is (1) $\beta_1 \phi(X\beta)$ when $x \leq \gamma$ and (2) $(\beta_1 + \beta_2)\phi(X\beta)$ when $x > \gamma$.

Figure 3: Contribution norms model for the extensive margin of punishment. Contributions below some contribution norm γ are more likely to be punished, while contributions above γ are less likely to be punished.

The problem with the linear probability model (besides generating probabilities below zero or above one) is that it predicts discrete jumps in the probability of punishment on either side of the norm. Panel 3b extends this to the nonlinear model, similar to the one used by Carpenter and Matthews (2009) and this paper, where Φ is the Normal CDF and the derivative ϕ is the Normal PDF. Now the probability of punishment is continuous at γ, ensuring smooth predictions around the norm. Though we cannot interpret γ as a hard threshold, the intuition remains: when stable

[9] Reuben and Riedl (2013) infer norms from punishments using a tobit model; we used the Carpenter and Matthews (2009) model because it separately estimates the probability of punishment and the magnitude of punishment, and we use these estimates to calculate the expected costs of noncompliance.

norms emerge, contributions below the norm are more likely to get punished, and contributions above the norm are less likely to get punished.

To estimate the norm γ in each treatment we ran a grid search on candidate models that varied the norm and estimated each model using maximum likelihood. The value of γ that maximized the likelihood is interpreted as the contribution norm. Carpenter and Matthews (2009) show that "absolute norms" (e.g., a subject must contribute a specific value of their choice set) outperform other candidates like the group average, so we only ran our grid search over the choice sets for *Low* and *High*.

Like Reuben and Riedl (2013), we collect the log-likelihood of each model and plot the normalized likelihood surface over all possible contribution norms (the worst-fitting norm is zero, and the best-fitting norm – the norm that maximizes the log-likelihood – is one). If γ is unique then we will see a single-peaked likelihood surface, where the peak is the norm. Otherwise the likelihood-surface will be relatively flat. We follow this procedure for both the extensive and intensive margins.

Lastly, to account for the period-by-period endogeneity of information in *Unobserved* we include a categorical variable in each regression that indicates the level of information subject i has in period t. We call this variable Reveal_t, and its values depend on i's endowment. If the subject is a *Low* type, then Reveal_t can be zero (no *High* types in their group contributed more than ten), one (*one* High type contributed more than ten), or two (both *High* types contributed more than ten). If the subject is a *High* type, Reveal_t simply indicates whether the other *High* type contributed above ten (in which case the variable is equal to one).

Our results proceed as follows for the both extensive and intensive margins. First, we present the likelihood surfaces and identify the emergent contribution norms. Second, we look at the average marginal effects of endogenous information. Since the rest of our coefficients are very similar to Carpenter and Matthews (2009), we relegate them to the appendices. Finally, we plug the contribution norms back into the models to calculate the expected costs of noncompliance.

3.2.1 Extensive margin of punishment

We estimated the probability of punishment using a random effects probit regression:

$$P(s>0)_{ijgt} = \Phi\Big(\alpha + \beta_1 x_{jgt} + \beta_2 \bar{x}_{gt} + \beta_3 (x_{jgt}-\gamma)^+ + \beta_4 \bar{x}_{gt}(x_{jgt}-\gamma)^+ + \mathbb{1}_{Unobserved}\beta_5 \text{Reveal}_t + \mathbf{Z'}_{igt}\psi + \mu_i + \varepsilon_{ijgt}\Big) \quad (2)$$

where $P(s>0)_{ijgt}$ is the probability that subject i punishes subject j in group g and period t, x_{jgt} is j's contribution in t, $\bar{x}_{gt}$ is the average contribution in group g in period t, $\mathbf{Z'}$ is a vector of controls including Period, i's contribution in t and i's received sanctions in $t-1$, μ_i is the random intercept for i, and ε_{ijgt} is the idiosyncratic error. Standard errors were clustered at the group level. The term $(x_{jgt}-\gamma)^+ = \max[x_{jgt}-\gamma, 0]$ describes j's deviation above the norm γ, and thus is turned on when $x_{jgt} > \gamma$. In other words, it is this term that allows target contributions to be treated differently on either side of the emergent norm γ.

Estimated contribution norms. We begin with *Observed*. We estimated separate models for each sender-type and target-type (e.g. Low→Low) using values for γ in the range of a target's endowment. So, *Low* targets were evaluated over the range $[0, 10]$, and *High* targets were evaluated over the range $[0, 30]$. The two figures on the left in Figure 4 show the likelihood-surfaces obtained from our grid search in *Observed*. *Low* senders are in blue, *High* senders in orange. The winning norm in each figure is marked with a dot.

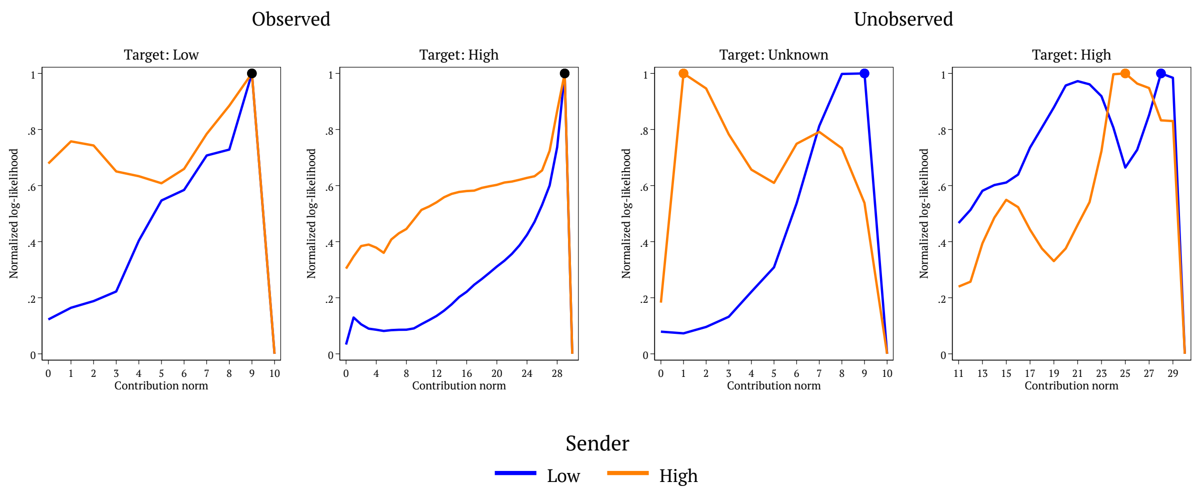

Figure 4: Likelihood surfaces of contribution norms for the extensive margin. The horizontal axis shows the candidate norm; the vertical axis shows the normalized log-likelihood of Equation 2 at that norm. Each graph shows the likelihood surface for *Low* senders (blue) and *High* senders (orange). The winning norm in each figure is marked by a dot.

Transparency leads to the emergence of efficient contribution norms. Both *Low* and High *were* expected to contribute their full endowments, as indicated by a norm of 9 for *Low* and 29 for *High*. In other words, *Low* types were expected to contribute *at least* 9, and *High* types were expected to contribute *at least* 29, so the probability of punishment was minimized when a subject contributed their full endowment. Moreover, there appears to be agreement over these norms. Both *Low* and *High* types enforce the same norms, and the likelihood surfaces are single-peaked.[10]

Next we turn to *Unobserved*. To account for period-by-period information we had to take a different approach in our grid search. Contributions over the range $[0, 10]$ pooled *High* and *Low* types, while contributions over the range $[11, 30]$ revealed *High* types. Therefore we estimated separated norms for the "unknown" range ($[0, 10]$) and the *High* range ($[11, 30]$). Combined with the information indicator, this allowed us to account for the information subject i had about target j in a given round. If j contributed between $[11, 30]$ then it was clear they were a *High* type. If j contributed between $[0, 10]$ while both *High* types contributed above ten, then it is clear that j

[10]Our results agree with Carpenter and Matthews (2009): the winning norm in their study, in which subjects had identical endowments of 25, was 24. In addition, our results are conceptually similar to Reuben and Riedl (2013). The authors use a different approach to identify contribution norms. Their free parameter, the corollary to our γ, describes how much a subject expected a target to contribute as a proportion of their own contribution. Like us, they find that subjects are expected to contribute their full endowment.

must be a *Low* type (the information indicator would equal two for *Low* and one for *High*).

The two figures on the right in Figure 4 show the grid search results for *Unobserved*. The first panel shows the results for the "unknown" range. The norm enforced by *Low* was 9, indicating that contributions that fell in this range were pushed up towards the *Low* endowment. *High* types enforced a norm of only one in this range, meaning they tended to start punishing at lower contributions.

Over the range $[11, 30]$ – when i knew for certain j was a *High* type – we see both *Low* and *High* enforce norms close to 30, at 28 and 25 respectively. In this sense we see consistency across treatments when subjects had complete information. A major difference across treatments in Figure 4 is that the likelihood surfaces in *Unobserved* are not single-peaked and show less agreement between *Low* and *High*. It appears that enforcing contribution norms was a messier affair in *Unobserved* compared to *Observed*.

Endogenous information. Recall that *High* types in *Unobserved* could reveal or conceal their endowments with their contributions in any period. In our regressions the variable Reveal_t captures this endogenous information. To distinguish levels of information (i.e. levels of Reveal_t) we name One High Reveal and Both High Reveal. These are calculated from the perspective of each sender – that is, they capture what each sender (*Low* or *High*) knew which contributions mapped to which types in a period. One High Reveal means that from the perspective of a sender, one of the three target contributions was greater than 10. To a *Low* type this means there was a fifty percent chance a contribution between zero and ten belongs to the other *Low* type. But to a *High* type, One High Reveal means they knew exactly who was *Low* or *High* (because they themselves are the other *High* type). Finally, Both High Reveal means that a *Low* type knew for sure that a contribution between zero and ten was from the other *Low* type.

The main result on endogenous information is the negative effect on Both High Reveal. The average marginal effect (AME) is significant for *Low* (AME $= -0.13, p = 0.08$). This means that when a *Low* type knew for certain they were facing another *Low* type, they were less likely to punish them. However, if the *Low* type were only fifty percent sure they were facing another *Low* type (the One High Reveal indicator), this effect goes away (AME $= 0.06, p = 0.42$).

High types were neither more or less likely to punish a target they knew was a *Low* type (AME $= -0.06, p = 0.15$). This is probably because *High* types withdrew from norm enforcement in *Unobserved*. Compared to *Observed* where *High* types registered 415 punishment events (about 15% of all observations), in *Unobserved* they registered just 244 punishment events (about 8% of all observations).

3.2.2 Intensive margin of punishment

Next we estimate norms on the severity of punishment. The norm on the intensive margin indicates where the severity of punishment changes: contributions below the norm receive harsher punishments, while contributions above the norm receive milder punishments.

We need to point out that the lion's share of punishments are zero. This means we are working with less data, so estimates of the intensive margin will be noisier (this was also the case in Carpenter and Matthews (2009)). Nevertheless, estimating the intensive margin is useful because we can combine it with the extensive margin to calculate expected costs of noncompliance.

Because sanctions were bounded below at zero and integer valued, we estimated the expected sanction size from subject i to subject j using a random effects Poisson regression:[11]

$$\mathbb{E}[s_{ijgt}|s_{ijgt} > 0] = \exp\Big(\alpha + \beta_1 x_{jgt} + \beta_2 \bar{x}_{gt} + \beta_3 (x_{jgt} - \psi)^+ + \beta_4 \bar{x}_{gt}(x_{jgt} - \psi)^+ + \mathbb{1}_{Unobserved}\beta_5 \text{Reveal}_t + \mathbf{Z'}_{igt}\psi + \mu_i + \varepsilon_{ijgt}\Big) \quad (3)$$

where ψ is the contribution norm and the other covariates are the same as Equation 2, including the information indicator in *Unobserved*.

Estimated contribution norms. Figure 5 shows the results of our grid search. Despite different amounts of data, norms on the intensive margin are fairly consistent with norms on the extensive margin. In *Observed*, both *Low* and *High* enforced norms at the endowments. So in *Observed*, punishment severity on *Low* and *High* only fell when each type contributed their full endowment. However, the likelihood surfaces are jagged, suggesting some disagreement over the winning norms.

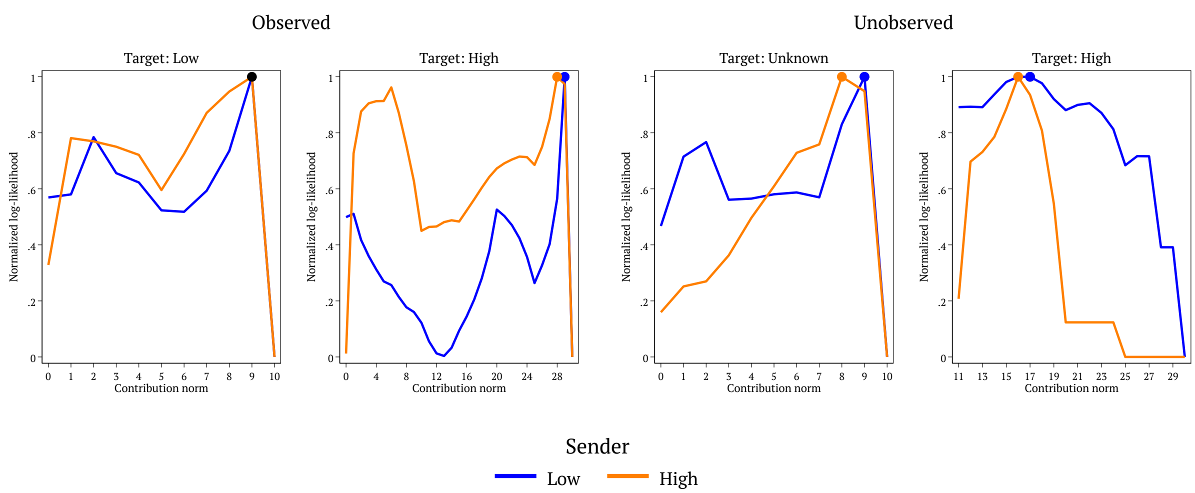

Figure 5: Likelihood surfaces of contribution norms for the intensive margin. The horizontal axis shows the candidate norm; the vertical axis shows the normalized log-likelihood of Equation 3 at that norm. Each graph shows the likelihood surface for *Low* senders (blue) and *High* senders (orange). The winning norm in each figure is marked by a dot.

For contributions less than or equal to ten in *Unobserved*, *Low* and *High* enforced a norm near the *Low* endowment. This is consistent with the extensive margin. However, when *High* types revealed (contributions between 11 and 30), both *Low* and *High* enforced lower norms of 16 and 17

[11] Carpenter and Matthews (2009) estimate the intensive margin using generalized least squares, but that may generate predicted sanctions below zero. Other studies use count data methods to estimate the intensive margin of punishment (e.g. De Geest and Stranlund, 2019).

respectively. This suggests that conditional on being punished, the norm that triggered more severe punishments to *High* types was smaller than in *Observed*.

Endogenous information. We find no significant AMEs on $Reveal_t$ on the intensive margin. For *Low* types the sign on Both High Reveal is the same as the extensive margin but the effect is insignificant (AME $= -0.41, p = 0.85$). The effect of One High Reveal is also negative and insignificant for *Low* (AME $= -0.27, p = 0.35$) and *High* (AME $= -0.19, p = 0.68$).

3.3 Expected costs of noncompliance

Overall, our results show that most punishment activity happened on the extensive margin. Interestingly, information did not dramatically effect contribution norms: across treatments the message was "contribute your endowment". But which endowment? Figure 2 shows *High* contributions peak at 10 (the *Low* endowment) and 30 (the *High* endowment). Next we look at how enforcement of norms led to the outcomes in *Observed* and *Unobserved.*

The costs of noncompliance increase (decrease) with larger deviations below (above) a norm. Since punishment was probabilistic, we trace out the expected costs of noncompliance by combining our estimates for the extensive and intensive margins.

We calculated the expected punishment C_{ij} from sender i to target j:

$$C_{ij} = P(s_{ij} > 0|x_j) \times \mathbb{E}[s_{ij}|s_{ij} > 0, x_j] \tag{4}$$

where $P(s_{ij} > 0|x_j)$ is the probability of punishment from i to j and $\mathbb{E}[s_{ij}|s_{ij} > 0]$ is the severity. Both terms were calculated by plugging in j's choice set $[0, e_j]$ into the derivatives of Equations 2 and 3 and evaluated with the estimated parameters at their likelihood-maximizing norms. Equation 3 ensures that predicted sanctions are bounded below at zero. To mimic the punishment technology in our design, we bounded predicted sanctions above at 10. We account for the distribution of types within groups (two *Low* and two *High*) when aggregating punishment. For instance, if target j is *High*, they can be targeted by the other *High* and two *Lows*, meaning total expected cost to j is the sum of punishments from two *Low* and one *High*.

We also accounted for the endogeneity of information in *Unobserved*. First we calculated the expected cost to *Low* and *High* in *Unobserved* when *Low* and *High* were pooled (both *Low* and at least one *High* contributed between zero and ten). Then we calculated the expected cost to *Low* and *High* when both *High* types reveal, implying that observed contributions between zero and ten came from a *Low* type.

Figure 6 shows our results. The horizontal axis shows the target's contribution, and the vertical axis shows the expected cost of punishment for that contribution in experimental dollars. In *Observed* (the solid lines in both panels), the expected costs line up pretty close to the deterrent level of punishment displayed in Figure 1. Contributions of zero by *Low* (*High*) were met with an expected cost of 6 (18). From those points the expected cost curves go down to about zero when both types contribute their full endowment. So in *Observed*, we not only see the emergence of efficient norms, we also see efficient norm enforcement.

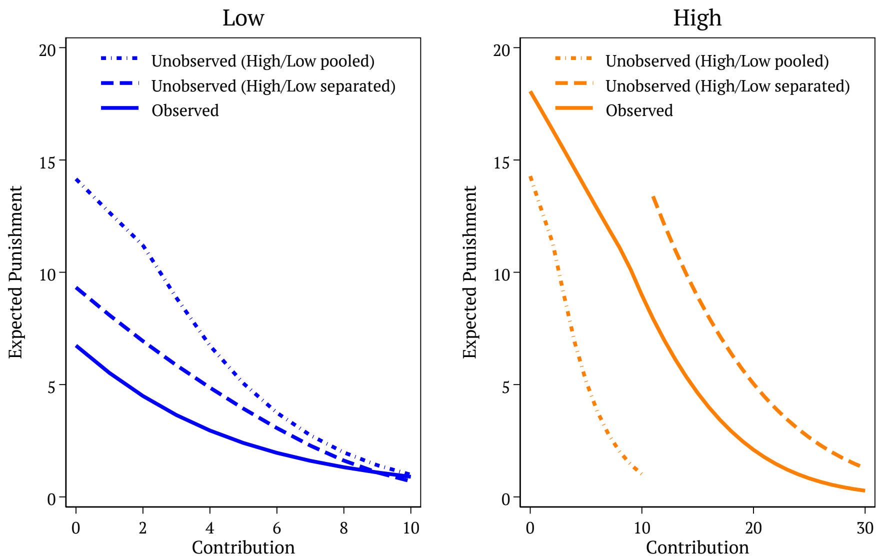

Figure 6: Expected cost of noncompliance for *Low* and *High* across treatments.

The picture changes in *Unobserved*. We calculated noncompliance costs according to the level of information available in a given period. The dot-dash line ("*High*/*Low* pooled") shows the expected costs to both *Low* and *High* when contributions were between zero and ten. The dashed lines ("*High*/*Low* separated") are expected costs when subjects in *Unobserved* had complete information in a given period (because both *High* types revealed).

The first thing to note is that *Low* types receive the same expected punishment when they contribute their entire endowment across treatments. However, it appears that deviations from full cooperation were treated more harshly in *Unobserved* relative to *Observed*. This would make sense if groups were reacting to the possible presence of *High* types among these contributions. Indeed, expected costs to *Low* fall when groups have complete information. However, these costs are still higher than the expected costs to *Low* in *Observed*.

Looking at *High* types in *Unobserved* we see a stepwise cost curve emerge. *High* types benefit from hiding, as contributions of 10 have expected punishments close to zero, compared to the expected punishments in *Observed* (the solid line). The moment a *High* type reveals (contributes above 10) the cost curve shoots up. But what is really interesting is how the cost curve for pooled contributions (the dot-dash line) seems to simply shift to the right for separated contributions (the dashed line). Costs are high at contributions of 0 and 11 but fall as contributions increase and settle close to zero at 10 and 30. This tells us that punishments were consistent: they either pushed subjects to contribute 10 or to contribute 30. Looking back at Figure 2 it is now straightforward to see from the expected cost curves why the distribution of *High* contributions peaked at 10 or 30. However, once again we see more punishment in *Unobserved*, as *High* contributions above 10 incurred larger costs than in *Observed*.

3.3.1 An evolutionary model of punishment with private information

Our empirical results show that the contribution norms that emerge under incomplete information *limit* how much agents with private information can hide, rather than prevent hiding at all. *High* types in *Unobserved* cannot get away with contributing nothing, but they can get away with contributing 10 and hiding behind a small endowment. It is possible that groups wanted to ensure some cooperation from *High* types while avoiding misguided punishment of *Low* types. This idea of agents trying to contain rather than eliminate the social costs of private information is also seen in bargaining games (Güth et al., 1996; Mitzkewitz and Nagel, 1993) and in public goods games with imperfect information (Nicklisch et al., 2016). To further motivate this idea we develop a simple evolutionary model of norm emergence based on our design.

Suppose two agents, one *Low* and one *High*, meet to play a public goods game with the same payoff function as our experiment, except we now set the MPCR to 0.8 (to ensure $\frac{1}{n} < \alpha < 1$) and restrict *Low* to two strategies (contribute 0 or 10) and *High* to three strategies (contribute 0, 10, or 30). As usual, the one-shot Nash equilibrium and subgame perfect equilibrium with a known end period is mutual defection (both contribute 0). Since our main interest is understanding the motives behind the emergence of the punishment rules in Figure 6, we will consider what happens when strategies in this game face different punishment rules imposed top-down and evolve according to a standard replicator dynamic seen in other public goods games (e.g. Cressman and Tao, 2014; Carpenter, 2004; Gintis et al., 2001; Miller and Andreoni, 1991).[12]

Suppose punishments are meted out by a social planner at zero cost. Since we have an equal population of *Low* and *High* types, we assume half the population of strategies are split among *Low's* two strategies and the other half are split among *High's* three strategies.

Figure 7 shows different simulations of this model under different punishment rules. *Low* is in blue and *High* is in orange, and the legend displays the population share of a strategy in the final time step. When there is no punishment, freeriding it sweeps through the population (Panel A). But if the social planner can observe the endowment of each type and apply exactly deterrent penalties according to the target's endowment (Panel B), cooperation emerges and both types contribute their full endowments.

[12]More details on the replicator dynamic in our model are in Section C of the appendix. Carpenter (2004) points out that replicator dynamics are a convenient way to mimic the learning process of groups in experiments.

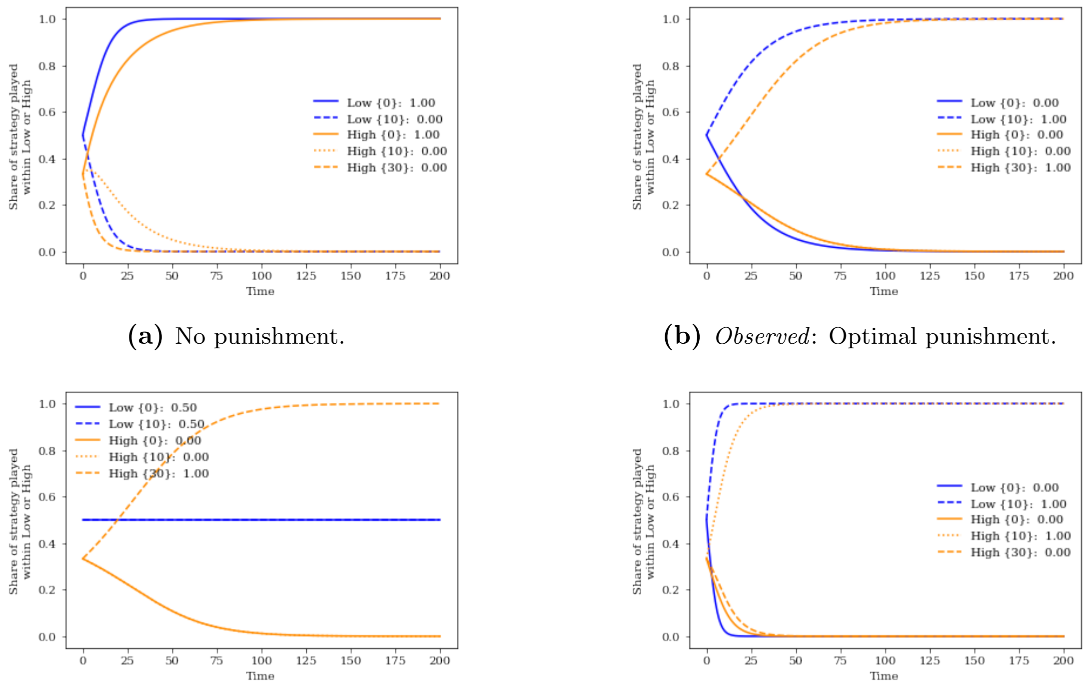

(a) No punishment.

(b) *Observed*: Optimal punishment.

(c) *Unobserved*: punish all 0 and 10 as *High*.

(d) *Unobserved*: punish all 0 as *High* and 10 as *Low*.

Figure 7: Replicator dynamics for a linear public good game with punishment, two endowment types and complete or incomplete information.

Now suppose the planner cannot observe endowments and must choose a punishment rule subject to this constraint. If the planner aims to secure full cooperation from *High*, they could punish all instances of 0 and 10 as if they came from *High*. Panel C shows the outcome of this scenario. Such a punishment rule secures full cooperation from *High*, but it reduces cooperation from *Low* by half and replaces it with *Low* freeriding, since both strategies earn the same payoff net of their respective punishments when *High* is fully cooperative.

Alternatively, the planner could instead adopt a punishment rule similar to what we see in our results: punish all contributions of zero as if they came from *High*, but do not punish contributions of 10 in order to avoid punishing cooperative *Low* types. This leads to Panel D, where, similar to our experimental results, we see full cooperation from *Low*, and partial cooperation by *High*.

So, it is plausible that subjects in our experiment evolved punishment rules that reflected a desire to protect *Low* rather than attack *High*. To some degree groups are cutting losses, since the rule allows *High* types to partially hide. However, this probably goes a long way towards preventing the unraveling of cooperation seen in games with imperfect information (e.g., Ambrus and Greiner, 2012).

4 Discussion and concluding remarks

We show that in a public goods game with peer punishment, heterogeneous endowments (*Low* and *High* types) and incomplete information (subjects observe contributions but not endowments), groups evolve contribution norms and penalties for noncompliance that: a) prevent *High* types from freeriding; b) reward *Low* types for cooperating; and c) enable *High* types to hide behind "small endowments", similar to how proposers in bargaining games with private information hide behind "small cakes" (Güth et al., 1996; Mitzkewitz and Nagel, 1993). These norms and their enforcement were effective because they prevented cooperation from completely unraveling. Groups selectively used punishment to attenuate noncompliance rather than eliminate it. A simple evolutionary model suggests that this approach is better than trying to prevent *High* types from hiding at all.

Our results compliment Reuben and Riedl (2013) and suggest the benefits to transparency in social dilemmas are twofold. First, transparency allows groups to map behavior to capacities (e.g., contributions to endowments), paving the way for efficient norms – norms that maximize aggregate payoffs – to emerge. Second, our estimates of the expected costs of noncompliance show that transparency allows groups to price discriminate when enforcing norm compliance and thus minimize the costs of enforcement. But when agents have private information and exploit it, we still see efficient norms emerge, but first-degree price discrimination (i.e. the cost of noncompliance depends of your endowment) is no longer possible. Moreover, our results show that in periods where groups in *Unobserved* had complete information, there was more expected punishment relative to *Observed*, suggesting that groups set too high a price of noncompliance.

In addition, our results clarify the effect punishment on cooperation with information asymmetries. Punishment with imperfect information (e.g., noisy signals about contributions) makes cooperation unravel (Ambrus and Greiner, 2012; Grechenig et al., 2010), and when given the choice, groups in noisy settings will discard punishment altogether (Nicklisch et al., 2016), largely to avoid misguided punishment of cooperators. Subjects in our experiment were also careful to avoid misguided punishment, but under incomplete information, this led to a stable level of cooperation. It would be interesting to see how a combination of imperfect and incomplete information might affect which norms emerge (if any) and their enforcement, since many real-world settings have both types of information asymmetries.

There are a number of ways to extend our paper and explore how and why norms emerge with private information. For starters, it is possible our results would change if we had larger groups or a different distribution of endowments. When *High* types pool, subjects during the punishment stage have to consider the probability a target is a *Low* or *High* type. In our design when both *High* types pool, the probability to *Low* that the target is *High* is $1/3$; when only one *High* pools, that probability increases to $1/2$. We show this has an effect on enforcement: *Low* types are more likely to punish a target when the probability they are *High* is $1/2$. Obviously these probabilities will change in larger groups or if there is a different distribution of endowments. Exactly how this might affect our results is unclear. Moreover, results could change if endowments are earned or randomly assigned: Jayadev and Bowles (2006) propose that spending on enforcement increases when inequalities are seen as illegitimate.

In addition, power asymmetries may play an important role. In our design we restricted *Low* and *High* types to the same enforcement budget each round. In reality, agents with more resources often have more power to influence outcomes at the macro-level (e.g. economic growth, Acemoglu et al., 2005) and the micro-level (e.g. the formation and enforcement of property rights Jayadev and Bowles, 2006). Power asymmetries may also influence institutional choice. Lower-endowment subjects vote for central enforcement over peer enforcement when there is incomplete information (De Geest and Stranlund, 2019), but those votes could be inconsequential if higher-endowment subjects have more power.

Another direction for future research is exploring whether private information leads to conflict over contribution norms. Nikiforakis et al. (2012) introduces the term "normative conflict" to describe settings with several plausible and appealing norms. The advantage of the linear game we used is that the Nash (contribute nothing) and social optimum (contribute everything) are easy for subjects to grasp, making it easier for them to impose and enforce efficient contribution norms under complete information. It also may help them adapt to the constraint of incomplete information, since the social optimum for the lower endowment subjects can serve as a focal point. However, this setting is much less realistic than a nonlinear game where normative conflict makes it harder for subjects to coordinate on what constitutes cooperation and non-cooperation (Kingsley, 2016; Cason and Gangadharan, 2015).[13]

Another question is the importance of enforcement. We show that norms emerge to attenuate rather than eliminate how much agents can exploit private information, and we show that the way subjects in *Unobserved* punished each other determined the emergence of these norms. However, direct enforcement could be less important to evolve and sustain norms if agents are already inclined to comply with norms.

We say this because we see somewhat similar results to ours in situations with incomplete information where noncompliance cannot be directly punished. For example, most cultures teach people not to lie, yet Abeler et al. (2019) show that in truth-telling experiments around the world, most subjects lie, but only a little. Abeler et al. (2019) argue this is because people like to be seen as honest[14] *and* because they have a preference for being honest. Another way to think about this is that subjects exploit their private information in the truth-telling task, but not as much as they could, because they are disciplined by the internalized norm "don't lie". If we are willing to speculate, the internalization of certain norms and the costs of noncompliance (e.g., guilt or some other psychological cost)[15] could lead to preferences for honesty, and perhaps even preferences for

[13]The interior design used in Kingsley (2016) introduces normative conflict along with incomplete information by creating two sets of contributions norms that both ensure efficiency (an equal contributions norm which maintains a payoff advantage for high endowment members or an equal payoffs norm which does not). Results suggest that low endowment members contribute an amount consistent with an efficient and equal payoff norm while high endowment members reject this norm and instead match the contributions of low members.

[14]A desire to be seen as honest (or fair) is also an argument used to explain behavior in other settings where a norm cannot be enforced like dictator games (Ockenfels and Werner, 2012; Andreoni and Bernheim, 2009) and charitable giving experiments (Grossman, 2015).

[15]Bucciol and Piovesan (2011) run a truth-telling task with children age 5 to 15 and find that a) children lie but less than they could, and b) children lie even less when the experimenter makes a normative appeal to tell the truth. Talwar et al. (2015) also show that normative appeals reduce lying among children in a non-incentized task. Interestingly, the authors also show that "expected punishments" (a child is told if they lie "you will be in trouble", but no punishment is actually carried out) crowd-out the effect of normative appeals and lead to an increase in lying.

norm compliance in general. Kimbrough and Vostroknutov (2016) make a similar point, suggesting that pro-social behavior can be explained by the fact that some subjects come into experiments with preferences to obey norms.[16] Kimbrough and Vostroknutov (2016) show that these "rule-followers" can sustain cooperation in a public goods game without punishment when paired with other "rule-followers", but they conditionally cooperate when paired with "rule-breakers" (i.e. "rule-followers" will not obey a norm if nobody else does). These arguments support the idea that enforcement drove norm emergence and compliance in our design, since subjects were not grouped together based on their inherent propensity to "do the right thing".

Putting these ideas together, it would be interesting to extend our design by combining multiple approaches to measuring norms and their compliance. For instance, the experimenter could start by measuring compliance propensity (Kimbrough and Vostroknutov, 2016) and explore different ways to match subjects along their propensities either exogenously (like Kimbrough and Vostroknutov, 2016) or endogenously (like Fehr and Williams, 2018). Next, elicit beliefs about norms within groups, perhaps with the coordination-game method of Krupka and Weber (2013) or the voting method of Fehr and Williams (2018). Finally, infer norms from punishment and estimate the expected costs of noncompliance.

[16] Kimbrough and Vostroknutov (2016) also point out that subjects often bring well-established norms into the lab (e.g. "do not lie", "split a surplus fifty-fifty") that regulate their behavior.

References

Abeler, J., D. Nosenzo, and C. Raymond (2019). Preferences for truth-telling. *Econometrica 87*(4), 1115–1153.

Acemoglu, D., S. Johnson, and J. A. Robinson (2005). Institutions as a fundamental cause of long-run growth. *Handbook of Economic Growth 1*, 385–472.

Ali, S. N. and D. A. Miller (2016). Ostracism and forgiveness. *American Economic Review 106*(8), 2329–48.

Ambrus, A. and B. Greiner (2012). Imperfect public monitoring with costly punishment: An experimental study. *The American Economic Review 102*(7), 3317–3332.

Andreoni, J. and B. D. Bernheim (2009). Social image and the 50–50 norm: A theoretical and experimental analysis of audience effects. *Econometrica 77*(5), 1607–1636.

Bornstein, G. and O. Weisel (2010). Punishment, cooperation, and cheater detection in "noisy" social exchange. *Games 1*, 18–33.

Bucciol, A. and M. Piovesan (2011). Luck or cheating? A field experiment on honesty with children. *Journal of Economic Psychology 32*(1), 73–78.

Carpenter, J. and P. H. Matthews (2009). What norms trigger punishment? *Experimental Economics 12*(3), 272–288.

Carpenter, J. P. (2004). When in Rome: Conformity and the provision of public goods. *The Journal of Socio-Economics 33*(4), 395–408.

Cason, T. N. and L. Gangadharan (2015). Promoting cooperation in nonlinear social dilemmas through peer punishment. *Experimental Economics 18*(1), 66–88.

Cressman, R. and Y. Tao (2014). The replicator equation and other game dynamics. *Proceedings of the National Academy of Sciences 111*(Supplement 3), 10810–10817.

De Geest, L. R. and D. C. Kingsley (2019). Endowment heterogeneity, incomplete information & institutional choice in public good experiments. *Journal of Behavioral and Experimental Economics*, 101478.

De Geest, L. R. and J. K. Stranlund (2019). Defending public goods and common-pool resources. *Journal of Behavioral and Experimental Economics 79*, 143–154.

Fehr, E. and T. Williams (2018). Social norms, endogenous sorting and the culture of cooperation. IZA Discussion Papers, No. 11457.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics 10*(2), 171–178.

Gächter, S., E. Renner, and M. Sefton (2008). The long-run benefits of punishment. *Science 322*(5907), 1510–1510.

Gintis, H., E. A. Smith, and S. Bowles (2001). Costly signaling and cooperation. *Journal of Theoretical Biology 213*(1), 103–119.

Grechenig, K., A. Nicklisch, and C. Thöni (2010). Punishment despite reasonable doubt? A public goods experiment with sanctions under uncertainty. *Journal of Empirical Legal Studies 7*(4), 847–867.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association 1*(1), 114–125.

Grossman, Z. (2015). Self-signaling and social-signaling in giving. *Journal of Economic Behavior & Organization 117*, 26–39.

Güth, W., S. Huck, and P. Ockenfels (1996). Two-level ultimatum bargaining with incomplete information: An experimental study. *The Economic Journal 106*(436), 593–604.

Jayadev, A. and S. Bowles (2006). Guard labor. *Journal of Development Economics 79*(2), 328–348.

Kimbrough, E. O. and A. Vostroknutov (2016). Norms make preferences social. *Journal of the European Economic Association 14*, 608–638.

Kimbrough, E. O. and A. Vostroknutov (2018). A portable method of eliciting respect for social norms. *Economics Letters 168*, 147–150.

Kingsley, D. C. (2016). Endowment heterogeneity and peer punishment in a public good experiment: Cooperation and normative conflict. *Journal of Behavioral and Experimental Economics 60*, 49–61.

Krupka, E. L. and R. A. Weber (2013). Identifying social norms using coordination games: Why does dictator game sharing vary? *Journal of the European Economic Association 11*(3), 495–524.

Miller, J. H. and J. Andreoni (1991). Can evolutionary dynamics explain free riding in experiments? *Economics Letters 36*(1), 9–15.

Mitzkewitz, M. and R. Nagel (1993). Experimental results on ultimatum games with incomplete information. *International Journal of Game Theory 22*(2), 171–198.

Nicklisch, A., K. Grechenig, and C. Thöni (2016). Information-sensitive Leviathans. *Journal of Public Economics 144*, 1–13.

Nicklisch, A. and I. Wolff (2011). Cooperation norms in multiple-stage punishment. *Journal of Public Economic Theory 13*(5), 791–827.

Nikiforakis, N., C. N. Noussair, and T. Wilkening (2012). Normative conflict and feuds: The limits of self-enforcement. *Journal of Public Economics 96*(9-10), 797–807.

Ockenfels, A. and P. Werner (2012). 'hiding behind a small cake' in a newspaper dictator game. *Journal of Economic Behavior & Organization 82*(1), 82–85.

Reuben, E. and A. Riedl (2013). Enforcement of contribution norms in public good games with heterogeneous populations. *Games and Economic Behavior 77*(1), 122–137.

Talwar, V., C. Arruda, and S. Yachison (2015). The effects of punishment and appeals for honesty on children's truth-telling behavior. *Journal of Experimental Child Psychology 130*, 209–217.

A Coefficient estimates: extensive margin of punishment

After the grid search we estimated the parameters in Equation 2 at the winning norms. Notice that the contribution norm γ appears twice in Equation 2. If punishment in our data mirrors Figure 3b, we should see $|\beta_1| > |\beta_1 + \beta_3|$, indicating that contributions below the norm are more likely to be punished, and contributions above the norm are less likely to be punished. However, Carpenter and Matthews (2009) find $\beta_3 > 0$, meaning contributions above the norm are *more* likely to be punished. While surprising, it is possible that such anti-social punishment is less likely to occur in more cooperative groups. This motivates the interaction term $\bar{x}_{gt}(x_{jgt} - \gamma)^+$, which allows the slope of the punishment curve after γ to depend not just on the target's contribution, but also on the overall cooperativeness of the group. Carpenter and Matthews (2009) find a significant and negative effect of this interaction; so, looking at Equation 2, we expect to see $\beta_4 < 0$.

Table 2 shows our results. We start with *Observed*. Across the board (and consistent with Carpenter and Matthews (2009)) the coefficient for Target Contribution is negative, meaning the probability of punishment decreased as the target's contribution approached the norm from below (corresponding to the segment labeled (1) of the curve in Figure 3b), although the effect is not significant for *High* types targeting *Low* or *High* group members.

The positive and significant coefficient on Average Contribution indicates that more cooperative groups were more likely to enforce norms. More cooperative groups were also less likely to engage in anti-social punishment. The positive coefficient on Deviation says that contributions at the norm (e.g. a *Low* type contribute 10 "deviates" above) were *more* likely to be punished. But this is offset by the negative interaction effect (Deviation $\times$ Average Contribution): in more cooperative groups, contributions at the norm (corresponding to the segment (2) of the curve in Figure 3b) were less likely to punished. There is some evidence of anti-social punishment: a *High* type who was punished in the previous round was more likely to punish a *Low* type in the next round.

Table 2: Estimated probability of punishment. Each model shows estimates for Equation 2 at the winning norms contribution norms (shown in Figure 4 and below each model number in the table). Models are organized by Sender $\rightarrow$ Target. For instance in *Observed*, $L \rightarrow L$ means "Low targeting Low"; in *Unobserved* $L \rightarrow [0, 10]$ means "Low targeting contributions between zero and ten".

	Observed				Unobserved			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	L $\rightarrow$ L	L $\rightarrow$ H	H $\rightarrow$ L	H $\rightarrow$ H	L $\rightarrow [0, 10]$	H $\rightarrow [0, 10]$	L $\rightarrow [11, 30]$	H $\rightarrow [11, 30]$
	$\hat{\gamma} = 9$	$\hat{\gamma} = 29$	$\hat{\gamma} = 9$	$\hat{\gamma} = 29$	$\hat{\gamma} = 9$	$\hat{\gamma} = 1$	$\hat{\gamma} = 28$	$\hat{\gamma} = 25$
Target Contribution	-0.462***	-0.163***	-0.131	-0.030	-0.188***	-1.429**	-0.146*	-0.168*
	(0.13)	(0.05)	(0.09)	(0.04)	(0.05)	(0.57)	(0.08)	(0.09)
Contribution	0.010	0.053***	-0.017	0.012	0.045	0.025	0.066**	-0.067
	(0.06)	(0.02)	(0.01)	(0.01)	(0.03)	(0.04)	(0.03)	(0.06)
Average Contribution	0.159**	0.130**	0.171**	0.057	0.117**	0.226**	0.265***	0.380**
	(0.07)	(0.06)	(0.08)	(0.06)	(0.06)	(0.11)	(0.08)	(0.17)
Lagged Sanctions	0.024	0.010	0.018**	0.012	0.017	0.032**	0.050***	0.020
	(0.03)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Period	-0.022**	-0.012	-0.007	-0.012	-0.006	-0.010	0.028*	0.017
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Deviation	0.666	0.800	0.525	1.829	2.349***	1.450**	1.433	–
	(0.96)	(1.22)	(0.55)	(1.44)	(0.37)	(0.64)	(1.14)	–
Average Contribution X Deviation	-0.163***	-0.211***	-0.117***	-0.175**	-0.177***	-0.023***	-0.125**	–
	(0.06)	(0.06)	(0.04)	(0.08)	(0.03)	(0.01)	(0.05)	–
One High Reveal					0.294	-0.575	–	–
					(0.38)	(0.42)	–	–
Both High Reveal					-0.986		-0.267	
					(0.61)		(0.29)	
Constant	-0.223	0.328	-1.887	-0.798	-1.152**	-1.275**	-2.538	-2.841***
	(0.82)	(0.79)	(1.18)	(0.89)	(0.54)	(0.60)	(1.90)	(0.92)
N	882	1764	1764	882	2232	2586	708	212
Log-likelihood	-79.687	-308.429	-348.752	-246.433	-825.153	-560.923	-167.181	-37.649

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered at the group level.

Many results in *Unobserved* are similar to those in *Observed* and suggest similar patterns in norm enforcement. Again the coefficients to Target Contribution are negative, the coefficients to Average Contribution are positive, and the coefficients to the interaction are negative, suggesting that more cooperative groups did more norm enforcement and less in anti-social punishment. We also see consistency among *Low* types across treatments: they who contributed more who were more likely to punish *High* types. In general, it appears that subjects in *Unobserved* enforced norms much as their counterparts did in *Observed*. At the same time, there is more evidence of retaliatory punishment in *Unobserved*. *Low* types were more likely to retaliate on *High* types. *High* types on the other hand were more likely to retaliate on "pooled" types (contributions between zero and ten).

B Coefficient estimates: intensive margin of punishment

Table 3 shows our parameter estimates for the winning norms. Generally, we find fewer significant results on the intensive margin, similar to Carpenter and Matthews (2009). This is likely due to the fact that there less data on the intensive margin. To take a stark example, *High* targeting *High* in *Unobserved* only has eleven observations.

Nevertheless, we do see consistency in the signs of parameter estimates on both margins. For instance, Target Contribution is negative, Deviation is positive, Average Contribution is positive, and the interaction Average Contribution X Deviation is negative. In addition, Both High Reveal is consistently negative, suggesting that punishment severity in general fell when groups had complete information.

Table 3: Estimated punishment given $P(Sanction = 1)$. Each model shows estimates for Equation 2 at the winning norms contribution norms (shown in Figure 4 and below each model number in the table). Models are organized by Sender → Target. For instance in *Observed*, $L \rightarrow L$ means "Low targeting Low"; in *Unobserved* $L \rightarrow [0, 10]$ means "Low targeting contributions between zero and ten".

	Observed				Unobserved			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	L → L	L → H	H → L	H → H	L → $[0, 10]$	H → $[0, 10]$	L → $[11, 30]$	H → $[11, 30]$
	$\hat{\gamma} = 9$	$\hat{\gamma} = 29$	$\hat{\gamma} = 9$	$\hat{\gamma} = 28$	$\hat{\gamma} = 9$	$\hat{\gamma} = 9$	$\hat{\gamma} = 16$	$\hat{\gamma} = 17$
Target Contribution	-0.108	-0.021**	-0.024	-0.089***	-0.132***	-0.139*	-0.057	-1.268***
	(0.23)	(0.01)	(0.22)	(0.03)	(0.02)	(0.08)	(0.47)	(0.14)
Contribution	0.050*	0.085	0.011	-0.037	-0.036	-0.004	-0.037	-0.150*
	(0.03)	(0.20)	(0.05)	(0.06)	(0.05)	(0.01)	(0.03)	(0.09)
Average Contribution	0.014	0.027	0.025	0.209	0.029	0.119**	0.124	1.232***
	(0.04)	(0.07)	(0.13)	(0.19)	(0.15)	(0.06)	(0.13)	(0.39)
Lagged Sanctions	-0.021	-0.017	0.002	-0.019	0.006	0.008	-0.003	-0.004
	(0.03)	(0.03)	(0.04)	(0.02)	(0.04)	(0.02)	(0.03)	(0.02)
Period	-0.014**	0.003	-0.001	0.001	0.007	0.016	0.011	-0.053***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.03)	(0.01)
Deviation	3.235	-0.488	1.533	1.230*	0.047	0.833	0.277	5.132***
	(3.65)	(0.47)	(2.53)	(0.71)	(1.28)	(0.73)	(0.51)	(0.45)
Average Contribution X Deviation	-0.145	-0.006	-0.073	-0.047	0.049	-0.062	-0.015***	-0.210***
	(0.12)	(0.03)	(0.21)	(0.04)	(0.08)	(0.06)	(0.00)	(0.02)
One High Reveal					-0.277	-0.194	–	–
					(0.29)	(0.47)	–	–
Both High Reveal					-0.413		-0.429	
					(2.21)		(0.28)	
Constant	0.421	0.031	0.300	0.375	1.259	0.320	0.161	5.537***
	(0.29)	(0.48)	(0.52)	(0.84)	(1.08)	(0.63)	(5.44)	(1.43)
N	45	234	240	158	380	221	104	11
Log-likelihood	-62.429	-375.007	-391.019	-294.722	-601.163	-392.457	-142.846	-14.109

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered at the group level.

C Replicator dynamics

Assuming a large, well-mixed population of a fixed size, the replicator dynamic describes how the proportion of the population playing a given strategy evolves from one period to the next based on the fitness or payoffs to that strategy. Consider for example a Low type who contributes 10. In a population of size N the fraction of Lows who contribute 10 is $L_{10} = \frac{N_{L_{10}}}{N}$. The fitness to L_{10} is then the sum of payoffs to playing 10 weighted by the share of each strategy in the population:

$$f_{L_{10}} = L_0\pi(10,0) + H_0\pi(10,0) + H_{10}\pi(10,10) + H_{30}\pi(10,30) \tag{C.1}$$

where $H_{10}\pi(10,10)$ is the weighted payoff to Low when they contribute 10 and High contributes 10, and so on. Average population fitness is just the sum of these fitnesses weighted by the proportion of agents playing any of the five strategies:

$$\bar{f} = L_{10}f_{L_{10}} + L_0 f_{L_0} + H_{10}f_{H_{10}} + H_{30}f_{H_{30}} + H_0 f_{H_0} \tag{C.2}$$

so the replicator dynamic for any strategy, for example Low playing 10, is then

$$\frac{dL_{10}}{dt} = \dot{L}_{10} = L_{10}(f_{L_{10}} - \bar{f}) \tag{C.3}$$

which simply says that the share of Low types contributing 10 will increase over time when the fitness of contributing 10 is greater than the average fitness, and will decrease if the opposite is the case.

D Experiment instructions

1 Welcome to the Experiment

Thank you for participating in our decision making experiment. The experiment consists of **50 periods**. In each period you will have an opportunity to earn money, which is in addition to the $5 guaranteed for your participation in the experiment. Your earnings each period will depend on your decisions and the decisions of other participants.

Please read the following instructions carefully. Everyone must correctly answer the comprehension questions at the end before we can begin.

During the experiment you are not allowed to communicate with other participants. If you have a question please raise your hand.

During the experiment your earnings will be calculated in *Experimental Dollars* (*EDs* for short). You can earn *EDs* every period. At the end of the experiment, your total earnings in *EDs* will be converted to U.S. dollars at the following rate:

$$100\,EDs = \$1$$

At the end of the experiment your total earnings (including the $5 participation payment) will be paid to you, privately and anonymously, in cash.

In the experiment, each participant is randomly assigned to a group of 4. This means that you are in a group with 3 other participants. You will be part of the **same group** throughout the entire experiment. However, at no point will the members of your group be revealed. All of the decisions you make within the experiment are anonymous and will be kept confidential.

In every period, each group member, yourself included, will be given an endowment of EDs. Two (2) members of the group will receive 30 EDs and two (2) members of the group will receive 10 EDs. This initial allocation of EDs is random and will be maintained throughout the experiment. Whatever your endowment is in Period 1 will remain your endowment for the entire experiment.

Each period consists of two stages. We will discuss both stages in detail, along with examples, and ask you to complete comprehension questions before starting the experiment.

Stage 1

Each of you will independently and anonymously decide how many of your EDs to allocate to the group account. You can allocate any integer between 0 and your endowment to the group account. Your remaining EDs will automatically be allocated to your private account. Your earnings depend on the number of EDs in your private account and the *total* number of EDs in the group account.

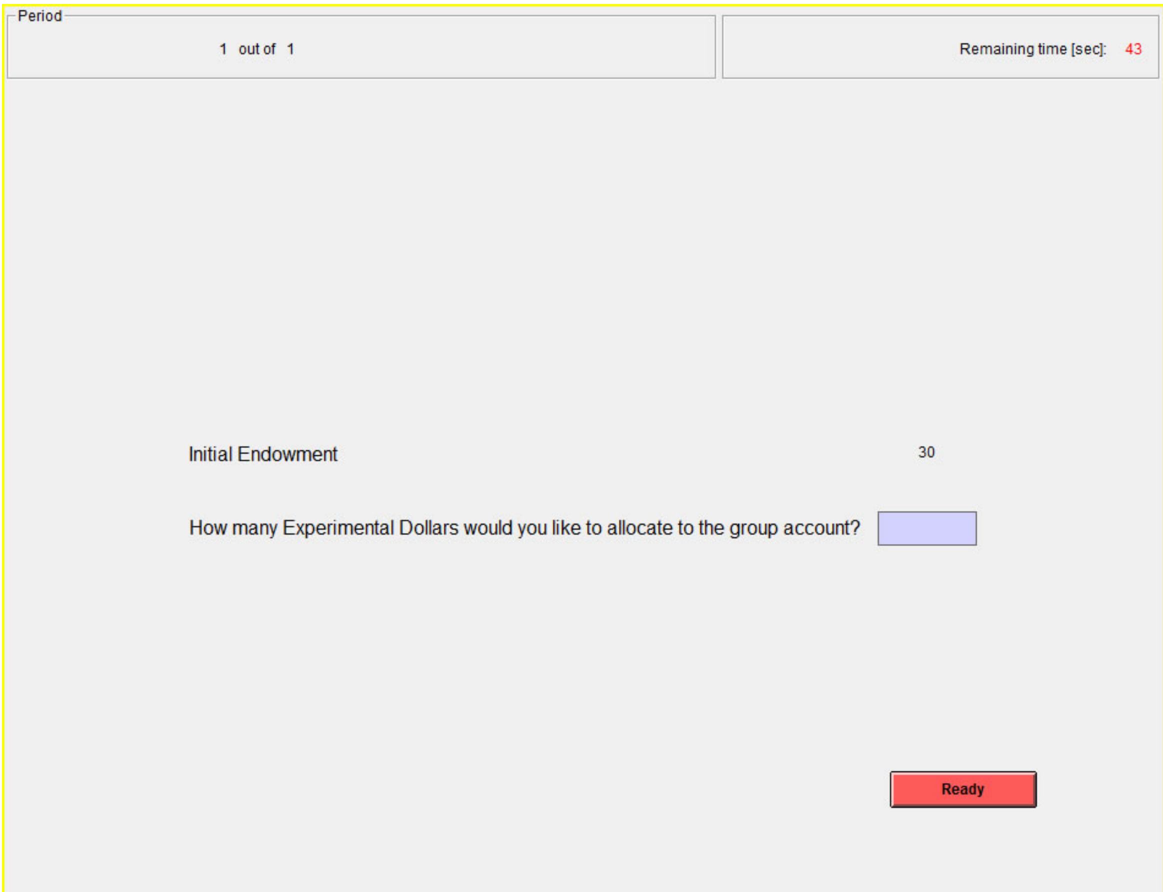

Figure 1: Example of allocation decision screen (assumes a 30 ED endowment)

How are period earnings calculated?

The earnings from your private account equal the number of EDs in your private account. Your private account earnings do not depend on the decisions of other group members. You simply keep all EDs that you choose *not* to allocate to the group account.

Your Private Account Earnings = (Your Endowment) - (Your allocation to group account)

Your earnings from the group account equal 0.4 times the *total* number of EDs allocated to the group account. Thus, your group account earnings depend, in part, on the decisions of other group members.

Your Group Account Earnings = 0.4*(the total number of EDs allocated to the group account)

Your period earnings are the sum of your private account earnings and your group account earnings.

Your Period Earnings = Your Private Account Earnings + Your Group Account Earnings

After Period 1 you will be presented with the history of your choices from previous periods. This information includes the information above and your total earnings up to this point in the experiment. Your total earnings are the sum of your earnings from each period of the experiment.

Your Total Earnings = Sum of your Private Earnings each Period

EXAMPLE 1

The example assumes the following:

	Endowment	Allocation to Group Account
You	30 EDs	15 EDs
Member A	30 EDs	30 EDs
Member B	10 EDs	10 EDs
Member C	10 EDs	0 EDs

The total number of EDs in the group account = 15+30+10+0 = 55 EDs, so each group member earns = 0.4*55 = 22 EDs from the group account.

What are your period earnings in this example?
You have a 30 ED endowment and you allocated 15 EDs:

> Your period earnings = private account earnings + group account earnings
> = (your endowment – your allocation) + (0.4*total group allocation)
> = (30 – 15) + 0.4*55
> = 15 + 22 = 37 EDs

What are the period earnings of Member A in this example?
Group member A has a 30 ED endowment and allocated 30 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (30 – 30) + 0.4*55
> = 0 + 22 = 22 EDs

What are the period earnings of Member B in this example?
Group member B has a 10 ED endowment and allocated 10 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (10 – 10) + 0.4*55
> = 0 + 22 = 22 EDs

What are the period earnings of Member C in this example?
Group member C has a 10 ED endowment and allocated 0 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (10 – 0) + 0.4*55
> = 10 + 22 = 32 EDs

Note that, regardless of your endowment, for each ED you allocate to the group account, your earnings from the group account ***increase*** by 0.4*1 = 0.4 EDs and your earnings from your private account ***decrease*** by 1 ED.

However, for each ED you allocate to the group account, the earnings of each of the other 3 members of your group ***increase*** by 0.4 EDs. Therefore, for each ED you allocate to the group account the total group earnings ***increase*** by 0.4*3 = 1.2 EDs.

You also obtain earnings from each ED allocated to the group account by others. You earn 0.4*1 = 0.4 EDs for each ED allocated to the group account by another member.

EXAMPLE 2

Relative to Example 1 assume that you decrease your allocation to 0 EDs but nothing else changes:

	Endowment	Allocation to Group Account
You	30 EDs	0 EDs
Member A	30 EDs	30 EDs
Member B	10 EDs	10 EDs
Member C	10 EDs	0 EDs

The total number of EDs in the group account = 0+30+10+0 = 40 EDs, so each group member earns = 0.4*40 = 16 EDs from the group account.

What are your period earnings in this example?
You have a 30 ED endowment and allocated 0 EDs:

> Your period earnings = (your endowment – your allocation) + (0.4*total group allocation)
> = (30 – 0) + 0.4*40
> = 30 + 16 = 46 EDs (An increase of 9 EDs relative to Example 1)

What are the period earnings of Member A in this example?
Group member A has a 30 ED endowment and allocated 10 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (30 – 30) + 0.4*40
> = 0 + 16 = 16 EDs (A decrease of 6 EDs relative to Example 1)

What are the period earnings of Member B in this example?
Group member B has a 10 ED endowment and allocated 10 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (10 – 10) + 0.4*40
> = 0 + 16 = 8 EDs (A decrease of 6 EDs relative to Example 1)

What are the period earnings of Member C in this example?
Group member C has a 10 ED endowment and allocated 0 EDs:

> Their period earnings = (their endowment – their allocation) + (0.4*total group allocation)
> = (10 – 0) + 0.4*40
> = 10 + 16 = 26 EDs (A decrease of 6 EDs relative to Example 1)

Compared with the earnings of Example 1, your earnings have increased, and the earnings of **each** of the other three members have decreased.

Comprehension 1

Please answer the following questions. Raise your hand if you need any help. A member of the experiment team will check your answers when you are done. We will begin when everyone has finished. Thank you for your patience.

1) Suppose that each group member, **including** you, allocates their *entire* endowment to the group account.

Suppose you have a 10 ED endowment and you allocate 10 EDs:

a What are your private account earnings? ________________

b What is the total number of EDs in the group account? ________________

c What are your group account earnings? ________________

d What are your period earnings? ________________

Now suppose you have a 30 ED endowment and you allocate 30 EDs:

a What are your private account earnings? ________________

b What is the total number of EDs in the group account? ________________

c What are your group account earnings? ________________

d What are your period earnings? ________________

2) Suppose that each group member, **including** you, allocates 0 EDs to the group account.

Suppose you have a **30 ED** endowment:

a What are your private account earnings? ________________

b What is the total number of EDs in the group account? ________________

c What are your group account earnings? ________________

d What are your period earnings? ________________

3) Suppose that each group member, **excluding** you, allocates 10 ED to the group account.

Suppose you have a 30 ED endowment and you allocate 0 EDs:

a What are your private account earnings? ________________

b What is the total number of EDs in the group account? ________________

c What are your group account earnings? ________________

d What are your period earnings? ________________

Assume you have a 30 ED endowment and you allocate 10 EDs:

a What are your private account earnings? ________________

b What is the total number of EDs in the group account? ________________

c What are your group account earnings? ________________

d What are your period earnings? ________________

STAGE 2

In each period, after Stage 1, your earnings are initially computed will be referred to as your ***Initial Period Earnings***. You will be shown:

- Your group account allocation
- The sum of the group account allocations by all members of your group
- Your group account earnings
- Your period earnings

In Stage 2, there will be a **deductions mechanism** which may affect your period earnings.

How does the deductions mechanism affect period earnings?

In each period, after each group member has made their allocation decision, each of you will continue to be shown the individual allocations and endowments of each group member by random ID.

Each group member will now have the opportunity to assign Reduction Points to other group members. The number of Reduction Points assigned can be any integer between 0 and 10 and can be distributed in any way among group members. Note that you don't need to assign any Reduction Points and you can only assign up to 10 Reduction Points. For each Reduction Point you assign to another group member you will pay 1 ED. This cost is referred to as:

Your Administrative Costs = The number of Reduction Points *you assign to others*

For each reduction point that is assigned to you your initial period earnings will be reduced by 3 EDs. This cost is referred to as:

Your Reduction Costs = 3 * The number of Reduction Points *assigned to you from others*

To calculate your period earnings you subtract your administrative costs and your reduction costs from your initial period earnings.

Note that your period earnings cannot be negative unless you assign Reduction Points. That is, you pay Administrative Costs.

Period Earnings = Max[Initial Period Earnings – Reduction Costs, 0] - Administrative Costs

Once each member has made their decisions concerning Reduction Points you will be shown:

- Your Administrative Costs
- Your Reduction Costs
- Your Period Earnings

1

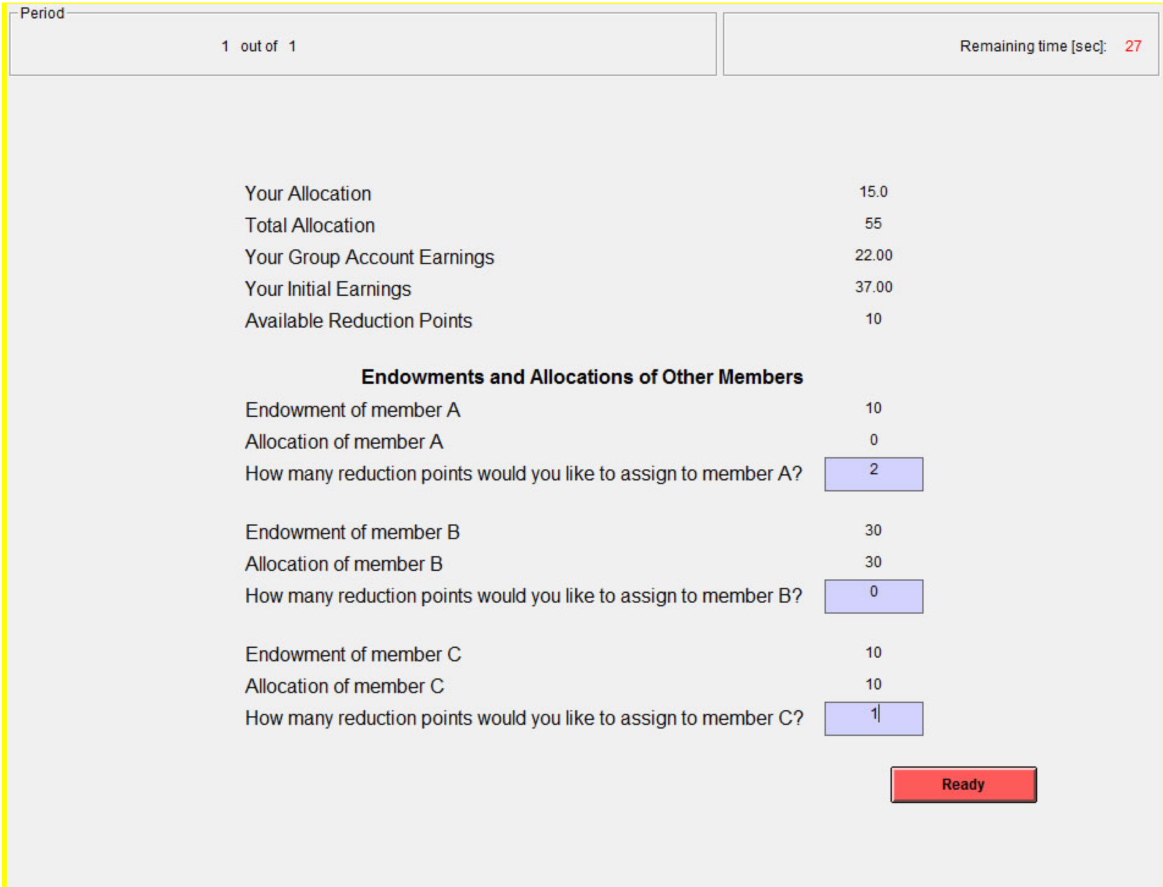

Figure 2: Example of your Reduction Point input screen given the example above.

EXAMPLE 3

The example assumes the following:

	Endowment	Allocation	Reduction Points Assigned	Reduction Points Received
You	30 EDs	15 EDs	2 to Member A 1 to Member C	1 from Member B
Member A	10 EDs	0 EDs	None	2 from you
Member B	30 EDs	30 EDs	1 to you	2 from Member C
Member C	10 EDs	10 EDs	2 to Member B	1 from you

Note that this information is provided for illustration only. You will not know how the other group members assigned their reduction points or which group members assigned reduction points to you (if any). In addition, you will not observe the endowments of other subjects, and they will not observe your endowment

The total number of EDs in the group account is = 15+0+30+10 = 55 EDs, so each group member earns = 0.4*55 = 22 EDs from the group account.

What are your period earnings in this example?
You have a 30 ED endowment, allocated 15 EDs, assigned 3 Reduction Points, and received 1 Reduction Points:

Your initial period earnings = private account earnings + group account earnings
= (your endowment – your allocation) + (0.4*total allocation)
= (30 – 15) + 0.4*55
= 15 + 22 = 37 EDs

Your administrative costs = 1 ED per Reduction Point you assigned (you assigned 3)
= 1*3 = 3 EDs

Your reduction costs = 3 EDs per Reduction Point assigned to you (you received 1)
= 1*3 = 3 EDs

Your period earnings = your initial period earnings – your administrative costs – reduction costs
= 37 – 3 – 3 = 31 EDs

1

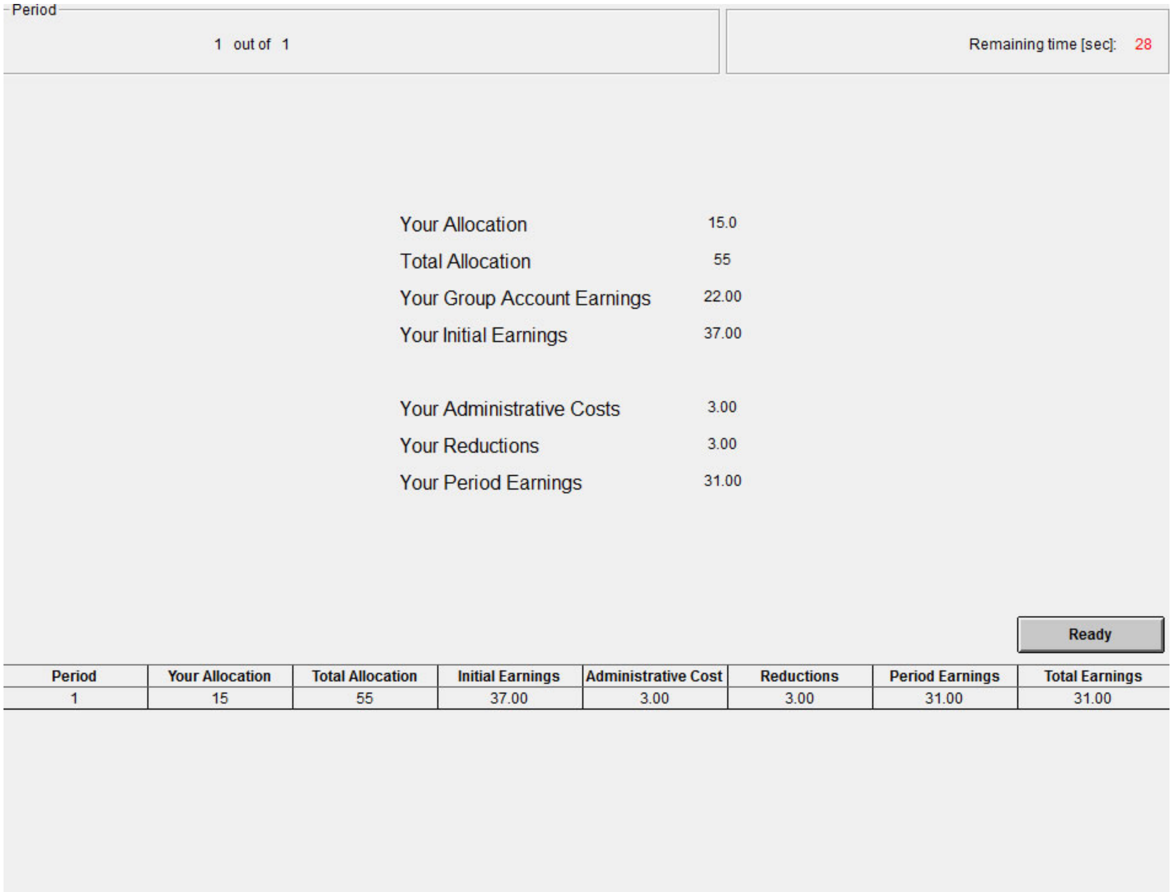

Figure 3: Example of your earnings screen given the example above

What are the period earnings of group member A in this example?
Member A has a 10 ED endowment, allocated 0 EDs, assigned 0 Reduction Points, and received 2 Reduction Points:

Member A's initial period earnings = (10 – 0) + 0.4*55
= 10 + 22 = 32 EDs

Member A's administrative costs = 1 ED per Reduction Point assigned (they assigned 0)
= 1*0 = 0 EDs

Member A's reduction costs = 3 EDs per Reduction Point received (they received 2)
= 2*3 = 6 EDs

Member A's period earnings = initial period earnings – administrative costs – reduction costs
= 32 – 0 – 6 = 26 EDs

Comprehension 2

Using the example above please answer the following questions. Raise your hand if you need any help. A member of the experiment team will check your answers when you are done. We will begin when everyone has finished. Thank you for your patience.

1. Determine the period earnings for Member B in the example above. Member B has a 30 ED endowment, allocated 30 EDs, assigned 0 Reduction Points, and received 3 Reduction Points.

   a What are Member B's initial period earning? ________________

   b What are Member B's administrative costs? ________________

   c What are Member B's reduction costs? ________________

   d What are Member B's period earnings? ________________

2. Determine the period earnings for Member C in the example above. Member C has a 10 ED endowment, allocated 0 EDs, assigned 2 Reduction Point, and received 2 Reduction Points.

   a What are Member C's initial period earning? ________________

   b What are Member C's administrative costs? ________________

   c What are Member C's reduction costs? ________________

   d What are Member C's period earnings? ________________