

Endogenous Group Formation via Unproductive Costs

JASON A. AIMONE

Virginia Tech Carilion Research Institute & Baylor University

LAURENCE R. IANNACCONE

Chapman University

MICHAEL D. MAKOWSKY

Johns Hopkins University

and

JARED RUBIN

Chapman University

First version received February 2011; final version accepted March 2013 (Eds.)

Sacrifice is widely believed to enhance cooperation in churches, communes, gangs, clans, military units, and many other groups. We find that sacrifice can also work in the lab, apart from special ideologies, identities, or interactions. Our subjects play a modified VCM game—one in which they can voluntarily join groups that provide *reduced* rates of return on private investment. This leads to both endogenous sorting (because free-riders tend to reject the reduced-rate option) and substitution (because reduced private productivity favours increased club involvement). Seemingly unproductive costs thus serve to screen out free-riders, attract conditional cooperators, boost club production, and increase member welfare. The sacrifice mechanism is simple and particularly useful where monitoring difficulties impede punishment, exclusion, fees, and other more standard solutions.

Key words: Endogenous group formation, Self-selection, Laboratory experiment, Free riding, Club goods, Voluntary contribution mechanism, Sacrifice, Religion

JEL Codes: C92, D71, H41, Z12

1. INTRODUCTION

Since the 1965 publications of Mancur Olson's *The Logic of Collective Action* and James Buchanan's *An Economic Theory of Clubs*, scholars have laboured to better understand how individuals can voluntarily organize to benefit from joint activities or shared resources. Central to this literature is the insight that success depends critically on a group's capacity to limit free riding, since the nature of club goods all but guarantees that individuals will be tempted to over-consume or under-contribute.¹

1. For clubs that exist to foster joint activity, the problem is not that people contribute nothing at all but rather that they contribute substantially less (or consume substantially more) than is collectively optimal. Cornes and Sandler (1984) label such problems "easy riding."

In theory, efficient clubs can be engineered through an appropriate combination of exclusion, membership fees, and usage tolls. In practice, the required mix often proves impossible, illegal, or excessively costly. Yet as Ostrom (1990) and other new institutional economists have shown, communities across the world and throughout history routinely craft ingenious solutions to seemingly intractable collective action problems.

This article provides the first experimental findings on a novel mechanism purported to reduce free-riding in real-world clubs. The mechanism is unusual in several respects. First, it is very simple. Second, it is seemingly counterproductive, as it rests on the voluntary *sacrifice* of valued resources. Third, it was first noted in research on religion. Fourth, both theory and case studies suggest that the mechanism works for a wide range of groups (including cults, sects, communes, clans, gangs, insurgents, fraternities, military elites, terrorists, ethnic subcultures, social movement organizations, and even some self-help groups). Fifth, the sacrifice mechanism does not require that individual inputs be accurately observed or appropriately rewarded. It can therefore remain viable where other methods fail—including usage tolls, graduated rewards, selective inclusion or expulsion, and most voluntary systems of arbitration, reputation, reciprocity, and punishment.

We study the impact of self-selected sacrifice within a standard laboratory game: the voluntary contribution mechanism (VCM) employed by hundreds of public goods experiments since the early 1980s. Our results demonstrate that sacrifice can operate as predicted, even when separated from the distinctive beliefs, identities, and interactions that characterize the cults, communes, and other real-world collectives claimed to enhance their solidarity through seemingly unproductive costs. Permitting subjects to endogenously sort into groups that provide *reduced* rates of return to private investment but unchanged rates of return to group contributions induces both screening and substitution.² The screening effect is pronounced. Among subjects classified as “free-riders” (through the method of Fischbacher *et al.*, 2001), fewer than one in five chose groups that required any sacrifice of private productivity. In contrast, more than half of all “conditional cooperators” and other non-free-riders chose groups that cut their private productivity by as much as 45%. Substitution occurred as the members of sacrifice groups increased their overall contributions—responding both to the decreased value of private investment and the anticipated increase in other sacrificers’ contributions. Endogenous sorting was key; a controlled experiment where subjects were exogenously sorted into groups of varying sacrifice levels found that price effects were substantially weaker than under endogenous sorting. In short, despite the presence of a few free-riders seeking to profit from the other sacrificers’ generosity, overall contributions in sacrifice groups substantially exceeded overall contributions in non-sacrifice groups, and contributors’ earnings were greater as well.

1.1. *Research on “sacrifice and stigma”*

The theory of “sacrifice and stigma” was first proposed as a rational-choice solution to the puzzle of seemingly gratuitous costs demanded by cults, communes, and other collectives. Iannaccone (1992, 1994) argued that prohibitions on otherwise normal behaviour—including restrictions on dress, diet, grooming, entertainment, economic activity, sexual conduct, and social interactions—could mitigate free-rider problems by raising the effective cost of activities that otherwise compete for members’ resources. Iannaccone showed that prohibitions could raise equilibrium utility in homogeneous populations through substitution alone. The potential gains

2. Our results are reliant on an exogenous sorting mechanism, but it is one that is known in advance and which allows subjects to *choose* their groups. The actual grouping, however, occurs through the mechanism.

were greater still in heterogeneous populations, where sacrifice could also serve to screen out the less-committed (Iannaccone 1992: 281–283, and 1994: 1205–1209).

Scholars in economics, sociology, anthropology, and political science have viewed numerous religious phenomena through the lens of sacrifice and stigma. They have invoked the theory to better account for the demographic characteristics of conservative denominations (Iannaccone, 1994), the commitment strategies of successful communes (Sosis, 2000), the high rates of giving among sect members (James and Sharpe, 2007), the global growth of Mormons, Jehovah's Witnesses, Evangelical Protestants, and Islam (Stark and Iannaccone, 1997; Chen, 2010; Hanson and Xiang, 2011), the appeal of Christianity over Greco-Roman paganism (Stark, 1996), the emergence of Reform and Ultra-Orthodox Judaism (Koyama and Carvalho, 2011), the decline of liberal American Protestantism (Finke and Stark, 1992), and much more. In other studies, sacrifice links a variety of religious and political-economic outcomes. Examples include Berman's (2000) study of rising fertility and declining employment among Ultra-Orthodox Israelis, Carvalho's (2013) model of veiling and the unintended consequences of veiling restrictions in contemporary Islam, Berndt's (2007) description of the success of Jews and other "high tension" religious minorities as middlemen, Richardson and McBride's (2009) theory of medieval craft guilds, Abramitzky's (2008, 2009) tests of labour migration theories with data on Israeli kibbutzim, Iannaccone and Berman's (2006) account of sects as key providers of social services where markets and governments fail, and Berman's (2009) related work on the efficacy of Hamas, Taliban, and other radical religious groups as joint suppliers of religious authority, social services, and terrorism.³

Because the sacrifice model makes no special assumptions regarding the substance of a club good, it is often invoked as a candidate explanation for secular examples of extremism and deviance—including demeaning initiations, dangerous rites of passage, painful rituals, and deviant styles of speech and dress. Researchers have argued that these and other costly demands enhance commitment and cooperation within social clubs, fraternities, gangs, mafias, military units, terrorists, clans, social movement organizations, protest groups, disadvantaged minorities, ethnic subcultures, secular communes, political parties, and self-help groups.⁴

Though case studies and survey research suggest that sacrifice and stigma may operate as advertised in many different settings, we cannot ignore alternative interpretations. Insiders will justify their seemingly strange ways in terms of time-honoured traditions, hidden truths, or directly productive behaviour. Critics will emphasize ignorance, indoctrination, coercion, or outright irrationality.⁵ Others will claim that most examples of seemingly bizarre beliefs and behaviour simply reflect atypical preferences. Still others ascribe cost-induced commitment to "cognitive dissonance," "effort justification," or other non-rational mechanisms (Aronson and Mills, 1959; Festinger and Carlsmith, 1959).

We have shifted the study of sacrifice from field to lab. Whatever results, we can be fairly sure it does *not* trace back to unusual ideologies, deviant tastes, intense interactions, outcast identities, or any other special feature of cults, communes, clans, and the like. As in standard

3. Yet another line of related research employs agent-based simulations, to model the dynamics of "religious markets" as heterogeneous, quasi-rational agents move from group to group, and the groups in turn grow, decline, or drift across a (low-to-high cost) "church-to-sect" spectrum. See, *e.g.*, Makowsky (2011, 2012), Montgomery (1996), and Chattoe (2006).

4. See, *e.g.*, Rapoport and Weiss (2002) on Jewish minorities, Bose (2010) on environmental "sects," Lembke (2012) on Alcoholics Anonymous, Shimizu (2011) on terrorist and military forces, D'Amico (2008) on tattoos and gang identity, Posner (2002) on social norms, Sosis *et al.* (2007) on ritual scars in tribal cultures, Thies (2000) on successful communes, Cimino (2011) on abusive hazing of newcomers, and Bram (2010) on the Chinese Communist Party.

5. Accusations along these lines were widespread responses to the rapid growth and aggressive proselytizing of Moonies, Krishnas, and other "cults" of the 1970s, and decades later to the 9/11 suicide bombings (Iannaccone, 2006).

VCM experiments, our subjects divide their initial endowments between their personal accounts and a group account, knowing that the group total will be multiplied and then distributed back evenly to all group members. We alter the usual procedure only by permitting people to select themselves into groups that provide reduced returns on the moneys retained in personal accounts.

1.2. *Recent VCM research*

The linear VCM is the simplest of all club goods experiments and the most thoroughly studied. Because subjects maximize their individual earnings by contributing nothing while maximizing group earnings by contributing everything, it provides a natural starting point for studies of cooperation versus free-riding. As with other experiments that pit self-interest against group-interest, subjects routinely contribute far more than the self-interested Nash equilibrium.⁶ But in repeated play cooperation decays rapidly and contributions approach zero.⁷

With the basic results from VCM experiments no longer in dispute, attention has shifted towards increasing cooperation or slowing decay. Sorting has received special attention, because we now know that most people are at least “conditionally cooperative” and much more likely to maintain cooperation when separated from the minority who consistently free ride. Hence, by sorting players according to the magnitude of their first-round VCM contributions, Gunnthorsdottir *et al.* (2007) were able to dramatically slow the decay rate of contributions in subsequent rounds. Gächter and Thöni (2005) likewise induced higher contributions and reduced decay through an exogenous, subject-blind sorting mechanism. Burlando and Guala (2005) increased contributions by typing and sorting subjects based on a variety of different methods. Swope (2002) was able to raise contributions (but unable to consistently improve overall welfare) by excluding individuals who failed to meet an exogenously set contribution threshold in otherwise standard VCMs. For other examples of increased cooperation through sorting, see Gunnthorsdottir *et al.* (2010a), Gunnthorsdottir *et al.* (2010b), and Grimm and Mengel (2009).

Closer in spirit to our own work are recent experiments that increase contribution rates through *endogenous* sorting. Though methods vary greatly from one study to the next, they nearly all involve repeated rounds of play, feedback on the past contributions of other players or groups, and repeated opportunities to alter group membership by bidding for new members, ejecting less cooperative members, moving to more desirable groups, etc. Thus, in Coricelli *et al.* (2004) players ranked each other based on past play histories and then bid for preferred partners. The subjects in Page *et al.* (2005) received the contribution histories of all players in all groups, and could then pay to rank one or more potential partners and be re-grouped through a matching algorithm. In Ahn *et al.* (2008) players learned about average rates of giving across all groups, could signal desire to switch groups, and were then admitted (and/or released) based on the votes of non-switching players. In Charness and Yang (2010) players received information on past plays and then could unilaterally exit groups or stay and vote to exclude others and/or merge groups.

Recent VCM research also shows that information on past behaviour can foster cooperation without sorting—if players can reward, punish, or condemn other players’ behaviour. See, *e.g.*, Masclet *et al.* (2003), Gächter and Thöni (2005), Noussair and Tucker (2005), Anderson and Putterman (2006), Xiao and Houser (2005), and Houser *et al.* (2008), but also note that punishment often backfires.

6. Throughout the article, we refer to the “self-interested” Nash equilibrium as one with purely selfish agents who, in the context of our experiment, are only concerned with maximizing their own payout.

7. For reviews of the public goods and VCM literature, see Ledyard (1995), Chan, Mestelman, and Muller (2008), Brandts and Schram (2008), and Chaudhuri (2011).

There is, of course, a certain irony in mechanisms that require subjects, or even experimenters, to know the contributions of other players. Free-riding scarcely exists in real-world clubs where individual inputs are plain to see. Insofar as free-riding persists, it is best tackled head on, through club theory's standard mix of membership fees, usage tolls, and size limits. Sacrifice theory was never proposed for such settings, including charities that primarily pool money for collective causes.

Sacrifice is reserved for settings where individual inputs *cannot* be easily observed or rewarded. Religious groups are a case in point, but so too are most communes, clans, fraternities, social clubs, gangs, mafias, military units, terrorist organizations, and social movement organizations. In all these groups and many more success hinges on effort, enthusiasm, faith, solidarity, commitment, and other inputs that are inherently difficult to measure and repay. The problem is compounded in proselytizing religions and other growth-oriented groups, where optimal outreach requires substantial subsidies to newcomers whose future contributions can scarcely be guessed (McBride, 2007). In such settings, sacrifice can provide a "second-best" solution to free-rider problems.⁸ As Iannaccone (1992: 275) notes, it often is relatively easy "to observe and penalize mere involvement in competing groups ... [or demand] some salient, stigmatizing behaviour that inhibits participation or reduces productivity in alternative contexts: shaved heads, pink robes, or an isolated location does the job quite effectively."

In short, this article has two aims: first, to carry a popular theory into the lab and subject it to a relatively clean test; and second, to extend experimental studies of collective action to include a novel mechanism that operates even when members' inputs are all but invisible.⁹ Our results demonstrate that even in the austere environment of lab-based VCMs, self-selected sacrifice can reduce free riding, enhance collective action, and raise individual welfare. Moreover, the observed gains involve both the screening and the substitution predicted by theory, as free-riders avoid high-cost groups, and as high-cost joiners increase their contributions.

2. THE EXPERIMENT

In the standard linear VCM game, each player in a G -member group receives an endowment, I , decides how much to allocate to his individual account versus a group account, and then receives a total payout of $a(I - g_i) + r(g_i + \sum g_{-i})$, where g_i denotes the player's group contribution and $\sum g_{-i}$ the total contributions of the other members. As long as $r < a < rG$, the dominant self-interested individual strategy is to contribute nothing to the group, leading to Nash equilibrium payouts of aI . But contributing everything yields Pareto-preferred payouts of $(rG)I$. In VCM experiments, initial contributions typically average half the endowment, but decay sets in with repeated play.

2.1. "Normal" versus "Sacrifice" rounds

All subjects in our experiment played one "Normal" and two "Sacrifice" rounds of the VCM game. In both versions of the game, players were placed in anonymous four-member groups,

8. For sacrifice to increase group utility, the inside activity must also be a sufficiently close substitute for the penalized outside activities and must generate sufficiently large spillover benefits for other members. Our VCM satisfies both conditions because the payoffs are linear and spillovers are 3/4ths the total benefit.

9. For two other experiments that successfully induce self-sorting *without* information on past play, see Orbell and Dawes (1993) and Bohnet and Kubler (2005). As in our experiment, these experiments work by letting players choose among different versions of the game—in Orbell and Dawes (1993) they can opt out of playing prisoner's dilemma games altogether, and in Bohnet and Kubler (2005) they can bid on the right to play a prisoner's dilemma version that makes cooperation less risky.

given a 10-token endowment to divide between their private and group accounts, and then paid “experimental dollars” according to the formula:

$$\prod (g_i, g_{-i}) = (1-s^*) \cdot (10-g_i) + (0.40) \cdot (g_i + \sum g_{-i}) \quad (1)$$

where g_i is the number of tokens that subject i contributed to the group account, $\sum g_{-i}$ is the total contributed by the other three members, and $(1-s^*)$ is the rate of return on private investment. In the Normal round $s^*=0$ and groups were based on random sorting. In each Sacrifice round, subjects chose a preferred rate $(1-s_i)$, which could vary from 0.55 to 0.95, and which led to the grouping described below. Since the marginal rate of return on private investment exceeded 0.40 in all versions of the game, the self-interested Nash equilibrium was always for each player to contribute nothing and earn $(1-s^*)10$, whereas the Pareto-preferred solution was for all to contribute fully and earn 16.

The critical difference between Normal and Sacrifice rounds is that Sacrifice rounds permitted subjects to indicate a preferred rate of sacrifice used both to sort them into their groups *and* to determine the rate of return that the members of their group receive on private investment. So, while private investment in the Normal VCM round always returned 1.00 E\$ per token, in Sacrifice rounds each subject chose a preferred private account return $(1-s_i)$ from the set $\{0.55, 0.60, 0.65, 0.70, 0.75, 0.80, 0.85, 0.90, 0.95\}$. We omitted 0.50 and 1.00 from the list of potential return rates simply because they stand out as natural “anchor points.”

Ideally, each subject was then grouped with players who chose the same sacrifice levels s_i . In practice, however, when working with 12 or 16 subjects per session, subjects could only count on being grouped with others who chose similar levels. As explained to all subjects, the grouping procedure was to order all $(1-s_i)$ choices, then group the four subjects choosing the highest rates, then group the four subjects choosing the next highest rates, and so on, while breaking all ties randomly. In each group, the rate of return on private investment is $(1-s^*)$, where s^* is the average sacrifice level (s_i) chosen by the members of that particular group, rounded to the nearest 0.05.

Once subjects were grouped and told their rate of return $(1-s^*)$, they each made their private and group investment decisions, $(10-g_i)$ and g_i , thereby earning individual payoffs according to equation (1) above.¹⁰

2.2. Typing subjects and running the VCMs

Because screening is an important feature of the sacrifice model, each of our VCM rounds incorporated the type elicitation method introduced by Fischbacher *et al.* (2001). As in FGF, our subjects made both “conditional” and “unconditional” contribution decisions. The “unconditional” decision was simply a standard VCM decision—how much to contribute absent any information about other members’ contributions. The “conditional” decision, however, was a contribution schedule indicating how much the subject chose to contribute as a function of the average unconditional contributions of the other members. Actual payouts were then made after randomly selecting one member from each group, taking the other members unconditional contributions, and then determining the selected member’s contribution based on his conditional contribution list. The FGF method elicits a great deal of information about players’ preferences and types, while minimally altering choices and outcomes (Fischbacher *et al.* 2012).¹¹

10. Note that if all subjects played their money-maximizing strategies and expected others to do the same, then everyone would choose the same maximal rate of return to private investment (0.95), everyone would be sorted randomly, everyone would make zero contributions, and everyone would earn 9.50 E\$.

11. The FGF method potentially imparts a small bias to choices, because each player now faces a $(G-1)/G$ probability that his unconditional contribution will enter the average that determines the G -th player’s conditional

Each subject participated in three one-shot VCM games, though they were not told in advance how many games they would play. We used two different orderings of Normal and Sacrifice VCM games. In the first treatment, subjects played the Normal round first and a Sacrifice round second. We call this the “Experienced” ordering, because it gave subjects experience with the standard VCM game before having to deal with sacrifice. In the second treatment, which we call the “Inexperienced” ordering, subjects completed a Sacrifice round first and a Normal round second.¹² We used both orderings to control for learning effects, which might otherwise be confounded with sacrifice effects if the Normal round always came first.¹³ In all treatments, the final round was a second Sacrifice VCM. Supplementary Appendix A provides more details, including subject instructions.

Before making their VCM game decisions, subjects read the instructions while an experimenter read the instructions out loud. Subjects answered graded quiz questions and received \$0.25 for each correct answer.¹⁴ Subjects were truthfully informed that they would play only once with their current group. In addition to their experiment earnings, subjects received a five-dollar show-up fee. Subjects averaged about two and a half hours in the laboratory and total pay of \$23.26.¹⁵ We conducted all sessions at the Interdisciplinary Center for Economic Science (ICES) at George Mason University using subjects randomly recruited from GMU’s student body.

Our basic experiment included 232 subjects—10 sessions with 16 subjects each and 6 sessions with 12 subjects. Half of our sessions followed the “Experienced” ordering (Normal, Sacrifice, Sacrifice) and half the “Inexperienced” ordering (Sacrifice, Normal, Sacrifice).¹⁶

We also ran control sessions involving another 40 subjects—one session with 16 subjects and two with 12 subjects. The Sacrifice rounds of our control sessions did *not* permit subjects to choose rates of sacrifice, but simply randomized them across groups of differing sacrifice rates.¹⁷ Table 1 summarizes the structure of all three treatments.¹⁸

3. THEORY AND PREDICTIONS

3.1. *Setup*

Although VCM experiments routinely yield more cooperation than game theory predicts, the predictions form a natural baseline from which to interpret observed outcomes. Moreover, as noted

contribution. This inflates the expected return to unconditional contributions for any player who expects the average of other players’ conditional schedules to be increasing. Any such effect, however, applied uniformly throughout our experiment because we used the FGF method in all rounds.

12. In the Experienced ordering, subjects read the additional instructions associated with the Sacrifice VCM after completing the Normal VCM round. In the Inexperienced ordering subjects read the additional instructions associated with the Sacrifice VCM immediately following the basic game instructions. We did this so that subjects approached the first VCM game as a one-shot game without the anticipation of the second round game.

13. It is possible that the Inexperienced ordering will have a potential for distorting in the Normal round, since subjects will have already played one round. This makes it more difficult to discern subjects’ types in the Normal round. We discuss this in the following sections and show that the results are comparable in both orderings.

14. Subjects in the Experienced ordering averaged 7.32 of 9 answers correct on their first try, and subjects in the Inexperienced ordering averaged 7.44.

15. Experiments were conducted by hand. The length, show-up bonuses, and payment levels in the experiment were pre-reviewed and approved by GMU’s Human Subjects Review Board and fit into the normal range of experiments conducted with this GMU subject pool.

16. Each ordering consisted of five sessions of 16 subjects and three sessions of 12 subjects.

17. By running all control sessions with the “Inexperienced” ordering, we were better able to assess the impact of (exogenously imposed) “sacrifice” absent both screening and learning.

18. The “Experienced” sessions had 16, 16, 12, 16, 12, 12, 16, and 16 subjects (in sessions with ID numbers 1, 3, 5, 7, 13, 15, 17, and 19, respectively); the “Experienced” sessions had 12, 16, 16, 16, 12, 12, 16, and 16 (in sessions #2, 4, 6, 8, 12, 14, 16, and 18); and the “Control” sessions had 16, 12, and 12 subjects (in sessions #9, 10, and 11).

TABLE 1
Summary of treatments

Round 1	Round 2	Round 3	Num. of sessions	Num. of subjects	Average earnings	Show-up payment
Ordering, VCM type, and sorting method						
“Experienced” Normal (random)	Sacrifice (endogenous)	Sacrifice (endogenous)	8	116	\$18.36	\$5.00
“Inexperienced” Sacrifice (endogenous)	Normal (random)	Sacrifice (endogenous)	8	116	\$18.42	\$5.00
“Control” Sacrifice (random)	Normal (random)	Sacrifice (random)	3	40	\$17.50	\$5.00
Total			19	272	\$18.26	\$5.00

above, cooperation routinely decays towards self-interested Nash equilibria in repeated, one-shot play. Before turning to our results, we therefore describe and analyse a formal game whose choices and payoffs mirror those of our experiment.¹⁹ The game is fully detailed in Supplementary Appendix B.

The game includes N players, each in a group of size G . These groups operate as classic clubs, with each player deriving utility from his own private and group inputs as well as the inputs of other players in his club.

We assume two types of players, L and H (for “low” and “high” cooperators).²⁰ Each player chooses her group contribution, g_i . The utility of L-types depends entirely on earnings, whereas the utility of an H-type has a component $u(g_i)$ that depends on her group contribution. We assume that $u(0)=0$, $u' > 0$, and $u'' < 0$. Players know their own types and know also that the population share of H-types is α . We make no particular assumptions about the source of this utility—whether from pure altruism, anticipated reciprocity, or “warm glow” (Andreoni, 1990).

As in the experiment, different groups provide different rates of return to private activity inputs. Any player may seek membership in any group, but all members of any given group receive the group rate of return, known also as the group’s level of “sacrifice.” Each player i chooses a desired group level of sacrifice, $s_i \in \{s_0 < s_1 < \dots < s_M\}$, knowing how this choice affects group sacrifice and composition. Specifically, each player knows that the G players with the largest s -choices will be grouped together and assigned a group level of sacrifice equal to their average s -choice, then the next G , and so on, with all ties broken randomly. Groups with the lowest level of sacrifice, s_0 , are more properly called “no-sacrifice” groups.

Let s^* denote the resulting sacrifice level in i ’s group. After learning s^* , subject i must decide how much wealth (normalized to 1) to contribute to the group, knowing that his earnings will be $(1 - s^*)(1 - g_i) + r(g_i + \sum_{j \neq i} g_j(s^*))$, where r denotes the return rate on group contributions, g_i denotes i ’s contributions, and $\sum_{j \neq i} g_j(s^*)$ denotes contributions from the other members of i ’s group.

19. It does not suffice simply to cite Iannaccone’s (1992) results because his model assumes a sacrifice continuum, numerous players, and nonlinear payoffs. These assumptions sidestep corner solutions (in which players devote nothing to club inputs) and the complexities that arise when the sacrifice choice of a single player induces discontinuous changes in the composition or sacrifice levels of existing groups.

20. We have developed more general versions of the model, incorporating numerous discrete types, a continuum of types, risk aversion, $u(\cdot)$ as a continuous random variable, and/or $u(\cdot)$ as a function of both g_i and g_{-i} . The models become less tractable without simplifying assumptions but otherwise extend our basic predictions.

Player i therefore chooses g_i in order to maximize expected utility, U_i :

$$U_i = (1 - s^*) (1 - g_i) + r \left(g_i + \sum_{j \neq i} g_j(s^*) \right) + u(g_i), \quad (2)$$

where the last term is omitted if i is an L. L-types are risk-neutral money maximizers, and H-types have quasi-linear utility. As in the experiment, the overall group return to giving rG always exceeds the return on private investment $(1 - s^*)$, which in turn always exceeds the individual player's return on giving r . Hence, the self-interested Nash equilibrium always entails universal free-riding and zero giving, whereas the social optimum requires that everyone contribute their entire endowment.

An L-type player will always choose $g_i = 0$, but an H type will choose $\max\{0, g_i^*\}$, where g_i^* is defined implicitly by the equation: $u'(g_i^*) = 1 - s^* - r$.

Players choose sacrifice levels that maximize their expected payout. Again, denote the chosen sacrifice level as s_i , which influences the group sacrifice level, s^* . Conditional upon s_i , expected utility is

$$E[U_i | s_i] = E \left[(1 - s^*) (1 - g_i(s^*)) + r \left(g_i(s^*) + \sum_{j \neq i} g_j(s^*) \right) + u(g_i(s^*)) | s_i \right] \quad (3)$$

There are two fundamental sources of uncertainty for each player. The first is that she does not know how much the other players will give to the group. The amount given is a function of the other subjects' types, which are unobservable. Secondly, the subject does not know the mapping of s_i (the sacrifice level chosen) onto s^* (the sacrifice level played), since this depends on the choices of the other subjects.

3.2. Solving the model

Three classes of equilibria exist, which we refer to as *pooling*, *perfect screening*, and *imperfect screening* equilibria. The three equilibria are defined as follows.

Definition 1. A *pooling equilibrium* exists when all subjects choose the same group s_j .

Definition 2. A *perfect screening equilibrium* exists when all L-type subjects choose the no-sacrifice group, s_0 , and all H-type subjects choose some group $s_j > s_0$.

Definition 3. An *imperfect screening equilibrium* exists when some L-type subjects choose group s_0 , some L-type subjects choose group $s_j > s_0$, and all H-type subjects choose group s_j .

The precise conditions under which each class of equilibria arise are derived in Supplementary Appendix B. Roughly speaking, pooling equilibria can arise at the highest sacrifice level (s_M) when price effects governing the responsiveness of H-type giving to changes in sacrifice are sufficiently elastic, making the high-sacrifice group payout large enough that no one has reason to switch to a less costly group. A pooling equilibrium can also arise in the no-sacrifice group (s_0) when price effects are sufficiently small. This results in such small group payouts at higher levels of sacrifice that no one has reason to give up the maximized private payouts in s_0 . A perfect screening equilibrium can emerge if price effects are neither too large nor too small (so H-types

prefer high-sacrifice groups, yet L-types gain nothing from joining them). Finally, an imperfect screening equilibrium occurs when price effects are such that if some but not all L's join the H's in high-sacrifice groups, the L's earnings are equal in both high- and no-sacrifice groups yet the H's never gain from joining L's in the no-sacrifice group.²¹

Because sacrifice increases the relative price of private production, it can screen free-riders from high-sacrifice groups while inducing the H-types to substitute towards more group investment. But screening never occurs in pooling equilibria, and since our experiments show no sign of pooling, we focus on the screening equilibria. Our testable predictions concern which types enter which groups and how they act once in them.

Absent pooling, only L's join the no-sacrifice groups, so contributions there must be zero. H's only possible reason to join a high-sacrifice group is to contribute along with fellow H's. Hence,

Prediction 1. (Price effects and screening) *In both perfect and imperfect screening equilibria, average group contributions must be higher in high-sacrifice groups (s_j) than in no-sacrifice groups (s_0).*

Is it possible to separate price effects from screening? In either screening equilibrium, a non-zero number of L-types choose no-sacrifice while all H-types choose higher sacrifice. We can type subjects based on their propensity to free-ride. Hence,

Prediction 2. (Existence of screening) *In either perfect or imperfect screening equilibria, subjects with preferences to free-ride choose no-sacrifice groups (s_0) with greater frequency than subjects with preferences for giving.*

We can push this prediction further and distinguish between perfect and imperfect screening equilibria. In a perfect screening equilibrium, free-riders *always* choose low sacrifice groups, whereas an imperfect screening equilibrium has some of them choosing higher sacrifice groups (in order to benefit from the generosity of non-free-riders). In the context of the model, imperfect screening occurs when L-types enter high-sacrifice groups primarily populated with H-types. We call the L-types who enter high-sacrifice groups “infiltrators.”

Definition 4. *An infiltrator is an L-type player who joins a high-sacrifice group.*

Infiltrators reduce the expected payout in sacrifice groups, and sufficiently many infiltrators will push the payout *below* that which free-riders receive with certainty in s_0 . In contrast, fewer infiltrators leaves the expected payout above this level. Hence, an imperfect screening equilibrium requires a certain (endogenously determined) *ratio* of L's to H's in high-sacrifice group.

Prediction 3. (Infiltrators) *In an imperfect screening equilibrium, the ratio of L's to H's in the high-sacrifice group is $\left(\frac{r(G-1)g^*(S_j)}{s_j-s_0} - 1\right)$. This ratio equalizes total earnings of L's in both high-sacrifice and no-sacrifice groups, s_j and s_0 .*

The same logic has welfare consequences. Since infiltration makes L-types indifferent between groups, and since H's contribute while L's do not, the H's in high-sacrifice groups must receive *less* overall income than the L's.

21. As we prove in Supplementary Appendix B, there are no screening equilibria in which H's choose lower sacrifice than L's or share the no-sacrifice groups with L's.

Prediction 4. *In an imperfect screening equilibrium, the average total payout (private + group) is higher in the no-sacrifice group (s_0) than in the high-sacrifice group (s_j).*

Finally, consider the question of screening versus price (or “substitution”) effects. The mere existence of screening (as in prediction 2) does not determine the extent to which the higher contributions of sacrifice groups derive from screening or substitution. But *random* sorting provides a test. H-types always give $g_i^*(s)$, which increases with the group sacrifice level, s . Hence, if both H’s and L’s were *randomly* sorted into groups, the average group payout would be $\alpha g_i^*(s)$. In either a perfect or imperfect screening equilibrium, the actual group payout must be *higher* than $\alpha g_i^*(s)$ in high-sacrifice groups because the actual share of H’s in those groups must exceed their overall population share (because only L’s populate the low sacrifice group). Hence, we can assess the effect of screening on payouts by comparing results from random versus endogenous sorting. If only price effects matter, then group payouts would be the same under both forms of sorting. But if high-sacrifice groups benefit from screening, the following prediction holds:

Prediction 5. (Effects of screening) *In both perfect and imperfect screening equilibria, endogenous sorting yields higher group contributions relative to random sorting in high-sacrifice groups (s_j) but lower group contributions in the no-sacrifice group (s_0).*

4. EXPERIMENT RESULTS

Five basic results merit emphasis from the outset, starting with the simple but striking fact that many subjects did choose sacrifice. Given the chance to join groups with reduced private productivity and no direct compensation, half of all subjects did so. This share scarcely changed in the second Sacrifice round. Second, the decision to sacrifice did not look like “noise” behaviour. More than four-fifths of (FGF-typed) free-riders rejected the option, whereas most conditional cooperators embraced it. Third, average contributions correlated strongly with levels of sacrifice all across the spectrum, ranging from fewer than two tokens per member in the no-sacrifice group to about nine tokens per member in the highest sacrifice group. Fourth, successful sorting accounts for most of this effect. The correlation between contributions and sacrifice was much weaker in the control sessions, where subjects were randomly assigned to groups across the sacrifice spectrum. Fifth and finally, despite the obvious losses associated with sacrifice, and despite the presence of some free-riding “infiltrators,” earnings in high-sacrifice groups averaged as much as non-sacrifice groups. In short, the mechanism seems to have worked as advertised.

We turn next to specific results, organizing the analysis around the equilibrium and comparative static predictions of our formal model.²² Begin by recalling that players in sacrifice rounds faced private rates of return between 0.95 and 0.55 E\$ per token. Groups with the 0.95 rate bore the lowest sacrifice and so received the highest return to private investment. We refer to these as “minimal sacrifice,” “lowest sacrifice,” or “no-sacrifice” groups. By way of comparison, we speak of all other groups as “high sacrifice” groups, or merely “sacrifice” groups. Occasionally we examine patterns across the entire “spectrum” of sacrifice. These terms aid the analysis but played no role in the actual experiment, since subjects only chose among different “rates of return” on private accounts.

22. We focus on predictions from the formal model, because other possible predictions, such as those regarding the degree of warm glow that subjects receive from giving to the group (as in Crumpler and Grossman, 2008), require relatively strong assumptions about subjects’ utility from giving.




FIGURE 1

Average unconditional contribution to group accounts in Sacrifice VCM rounds

4.1. Unconditional contributions

Consider first the unconditional contributions in the no-sacrifice groups. We expect screening and price effects to induce higher average contributions in high-sacrifice groups than no-sacrifice groups (prediction #1). This indeed is what we find. As seen in Figure 1, sacrifice group players consistently averaged more than *twice* the contributions of their no-sacrifice group counterparts—whether playing the first or second sacrifice round, and whether having first played a Normal VCM or Sacrifice VCM.^{23,24}

As for variation across the entire sacrifice spectrum, prior studies (based on simulations, surveys, and formal theory) all claim that more sacrifice leads to greater giving. We see exactly this in Figure 2, which shows the dramatically lower average unconditional contributions in groups of lower sacrifice (or equivalent, higher private return rates, $1 - s$). The Cuzick trend test is statistically significant (at $p < 0.001$) in the first round of sacrifice (for Inexperienced alone, Experienced alone, and all subjects pooled) and for the second round as well.

23. We classify contribution results by group sacrifice level because subjects made their contribution choices only *after* learning their group assignments. The results are quite similar when reclassified according to subjects' preferred level of sacrifice.

24. Each of the sacrifice versus non-sacrifice differentials (\$3.64 versus \$1.78, \$5.61 versus \$2.52, and \$4.67 versus \$1.36) are statistically significant, with p -values under 0.001 for the two-sided Mann–Whitney test. Mann–Whitney tests assume equal variances between comparison groups. We therefore report any variances that differ at the 5% significance level. Variance of contributions in second round sacrifice groups is greater than in non-sacrifice groups (3.79 versus 2.39, two-tailed F -test, $p < 0.001$), though a t -test similarly shows a significant difference in contribution means, $p < 0.001$.




FIGURE 2

Average unconditional contribution to group account

4.2. Conditional contributions

Prediction 2 states that in both perfect and imperfect screening equilibria, free-riders choose non-sacrifice groups more frequently than subjects with preferences for giving. We test this prediction with our data on player types.

Specifically, for each of the three rounds of a session, we typed all subjects from their conditional contribution schedules, following the FGF method (Fischbacher *et al.*, 2001). We classified as “conditional cooperators” all subjects whose contribution schedules increased monotonically (or had a positive Spearman rank correlation coefficient with $p < 0.01$). The schedules of “free-riders” ran zero straight across (with the possible exception of one token at some single point). ‘Hump-shaped’ schedules rose and then fell. ‘Other’ schedules fit none of these patterns.²⁵

Figure 3 plots the resulting average contribution schedules in the Normal VCM round. The classification scheme yields 48% conditional cooperators, 23% free-riders, 15% ‘hump-shaped’ contributors, and 14% other. The pooled Sacrifice rounds yield a similar distribution of types with 52% conditional cooperators, 21% free-riders, 13% ‘hump-shaped’ contributors, and 14% others. These distributions are remarkably similar to those obtained by Fischbacher *et al.* (2001) and Fischbacher *et al.* (2012).²⁶

25. Supplementary Materials indicate precisely how each subject was typed.

26. In their standard VCM games, Fischbacher *et al.* (2001) classified 50%, 30%, 14%, and 6% of their subjects as conditional cooperators, free-riders, hump-shaped contributors, and other types, respectively. The corresponding shares in Fischbacher *et al.* (2012) were 55%, 23%, 12%, and 10%.

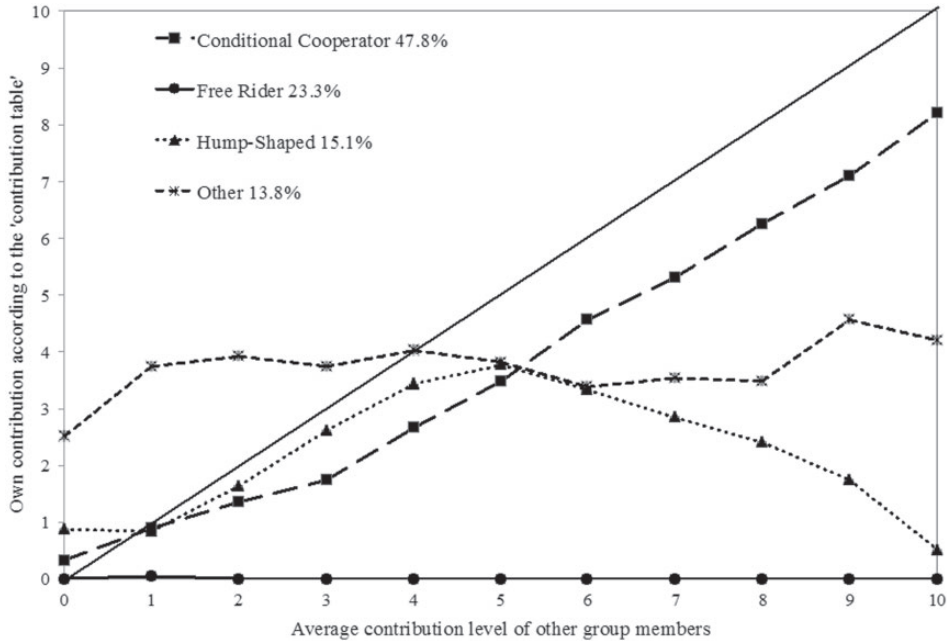


FIGURE 3
Subject “Types” and conditional contributions

Figure 4 confirms again that sacrifice tends to screen out free-riders. In both first and second sacrifice rounds, fewer than one in five “free-rider” types choose sacrifice groups: just 18% overall, compared to 53% for other types.^{27, 28}

Yet more evidence of successful screening comes from Experienced subjects. These subjects can be typed from their Normal VCM responses, which preceded any exposure to sacrifice and were independent of the sacrifice mechanism.²⁹ Only one of the 20 resulting “free-riders” subsequently chose non-zero sacrifice in the first Sacrifice round. In contrast, nearly half of the other types chose positive levels of sacrifice: 27 of 55 conditional cooperators, 10 of 23 hump-shaped, and 8 of 18 others.

The net effect of all this self-screening is seen in Figure 5, which contrasts the resulting composition of sacrifice versus non-sacrifice groups. As one might expect, the results strongly support the prediction (#2) that sacrifice groups end up with fewer free-riders and more conditional cooperators. Again, the difference is substantial. In the first Sacrifice round, Free-riders account for 32% of Non-Sacrifice group membership but only 10% of Sacrifice group membership.³⁰ In

27. These numbers reflect the private return rates that subjects actually *chose*, as opposed to the group rates they ultimately played. In some sessions, the grouping method (required to maintain constant group size) had the effect of allocating some would be 0.95 players to a 0.90 group (and in other sessions, vice versa).

28. This difference is statistically significant in each Sacrifice round, with $p < 0.001$ on the Mann–Whitney two-tailed test. The difference is similarly significant both for Experienced and Inexperienced subjects separately. Significance carries over to unequal variance t -tests.

29. Typing based on this round avoids endogeneity problems that might otherwise occur if the sacrifice level required by one’s group affected the overall shape of one’s conditional contributions.

30. The Mann–Whitney two-tailed test yields a $p < 0.001$. Significance holds also for Experienced ($p < 0.001$) and Inexperienced ($p = 0.017$) groups separately. T -tests likewise yield significant differences in mean frequency of free-riders although the mean variance in Non-Sacrifice groups is significantly greater than in Sacrifice groups.

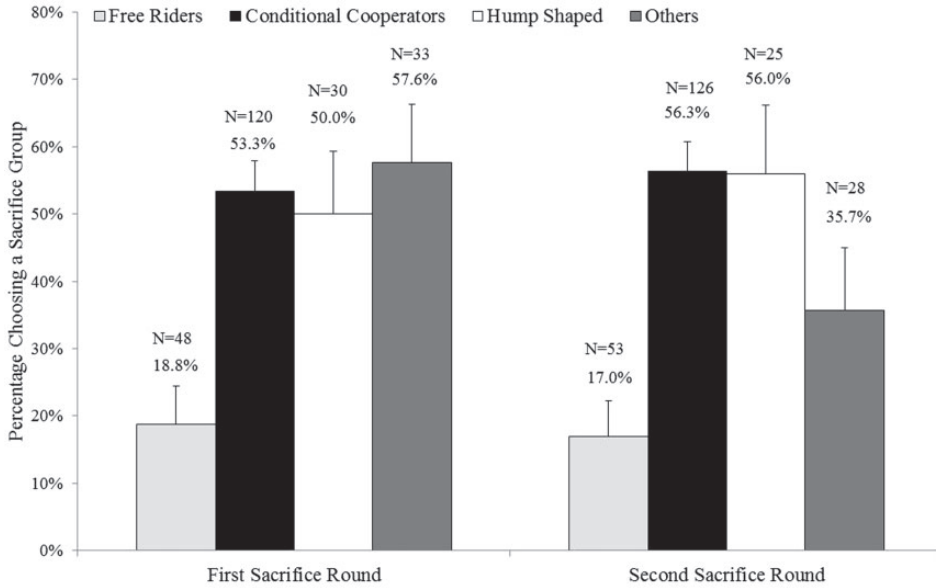


FIGURE 4
Percentage of each type choosing a sacrifice group

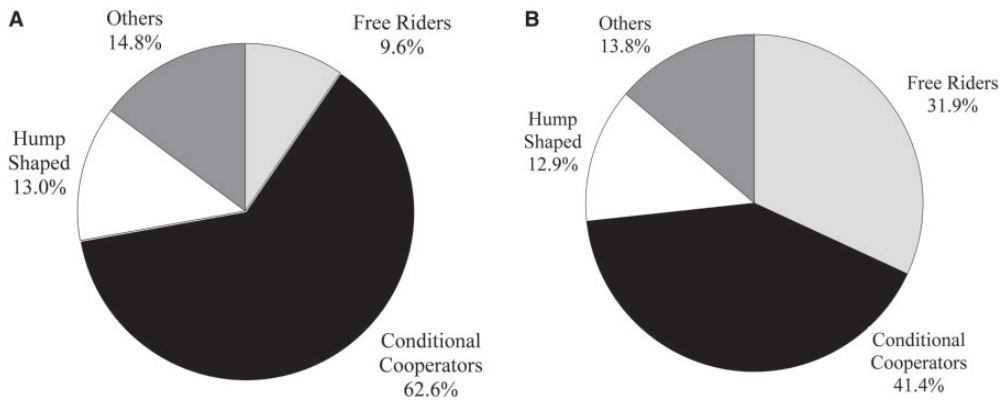


FIGURE 5
Composition of groups in first sacrifice round. (A) Sacrifice groups; (B) Non-sacrifice groups

contrast, conditional cooperators account for 63% of Sacrifice group membership but only 41% of Non-Sacrifice group membership.³¹

Figures 4 and 5 are consistent with an *imperfect* screening equilibrium, in which a non-trivial number of free-riders infiltrate the sacrifice groups (prediction 3).

“Infiltrators” enter high-sacrifice groups with no intention of contributing, but hoping that high contributions from others will more than offset their own reduced private earnings. Table 2

31. The *p*-values associated with this difference are less than 0.001 for all subjects combined, less than 0.01 for Experienced subjects alone, and *p* = 0.12 for Inexperienced subjects alone.

TABLE 2
Actual and predicted infiltrators and free-riders

		Sacrifice groups									Non-sacrifice group	
Private return =		0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	Total	Private return = 0.95	
First sacrifice round	Number of infiltrators	0	2	0	3	5	0	4	11	25	Free riders	56
	Expected infiltrators	3.4	6.7	3.7	3.4	7	-0.5	3.1	11.3	38.0		
	<i>Difference</i>	3.4	4.7	3.7	0.4	2.0	-0.5	-0.9	0.3	13.0		
	Total N	8	16	8	16	20	4	16	28	116	Total N	116
	Percentage infiltrators (%)	0.0	12.5	0.0	18.8	25.0	0.0	25.0	39.3	21.6	Percentage free-riders	48.3
Private return =		0.55	0.60	0.65	0.70	0.75	0.80	0.85	0.90	Total	Private return = 0.95	
Second sacrifice round	Number of infiltrators	1	3	4	1	1	5	3	6	24	Free riders	75
	Expected infiltrators	8.1	7.2	9.3	-1.2	2.1	4.2	8.7	2.1	40.6		
	<i>Difference</i>	7.1	4.2	5.3	-2.2	1.1	-0.8	5.7	-3.9	16.6		
	Total N	16	16	20	4	8	12	20	12	108	Total N	124
	Percentage infiltrators (%)	6.3	18.8	20.0	25.0	12.5	41.7	15.0	50.0	22.2	Percentage free-riders	60.5

Note: This table counts as “infiltrators” and “free riders” all subjects who contributed nothing to their group.

further confirms their presence. This table reports the number of infiltrators against the number of expected infiltrators predicted by the indifference condition in Prediction 3.³² For example, in the group with private return equaling 0.60 in the first sacrifice round, the indifference condition predicts that 6.7 of the 16 subjects should be infiltrators, whereas only 2 were in the experiment.

Recall that an imperfect screening equilibrium includes sufficient infiltration to make free-riders indifferent between sacrifice and non-sacrifice groups. As it turns out, actual infiltrators are about one-third fewer than predicted: 49 versus 79.

Fewer infiltrators means greater welfare for other members. Prediction 4 implies that imperfect screening leads to reduced overall earnings in high-sacrifice groups relative to low sacrifice groups. But in the experiment, infiltrators are so few that high-sacrifice groups earn slightly more.^{33,34} Figure 6 displays the average overall payouts to members of different groups across

32. Prediction 3 suggests that there should be $-1 + [0.12g^*(s_j)]/(s_j - s_0)$ infiltrators in group s_j . We derive the values in Table 2 by taking the average amount given in each group as the relevant $g^*(s_j)$. The formula in prediction 3 assumes zero contributions in no-sacrifice groups. In the experiment, however, some non-free-riders chose no-sacrifice groups, thereby yielding positive average group payouts. The calculations of Table 2 adjust for this added return.

33. Numerous phenomena can account for the fact that we see fewer infiltrators than expected and that the total payout is increasing in the sacrifice level. For one, agents in the model are risk neutral. Yet, if subjects are risk averse, then they are willing to pay a risk premium to stay in the non-sacrifice group, which is less risky. Risk aversion reduces the number of expected infiltrators while increasing the total payout of the high-sacrifice group. In fact, if agents were modelled as having a logarithmic utility function, $u(c) = \ln(c)$, then free-riders in the non-sacrifice group would require a premium of just 8.6% on their earnings (1.363E\$) to infiltrate at the rate that we find. Another possibility is that behavioural phenomena may be at work. For example, Iriberry and Rey-Biel (2009) find evidence that selfish subjects cannot imagine the existence of non-selfish subjects. If this is true to some degree among our subjects, then free-riders would underestimate the potential gains from infiltrating, which is in line with our findings.

34. We do not wish to make a welfare argument here. A simple extension to the model in which H-types receive more utility the more that others give to the group would entail that the sacrifice mechanism is welfare-enhancing for

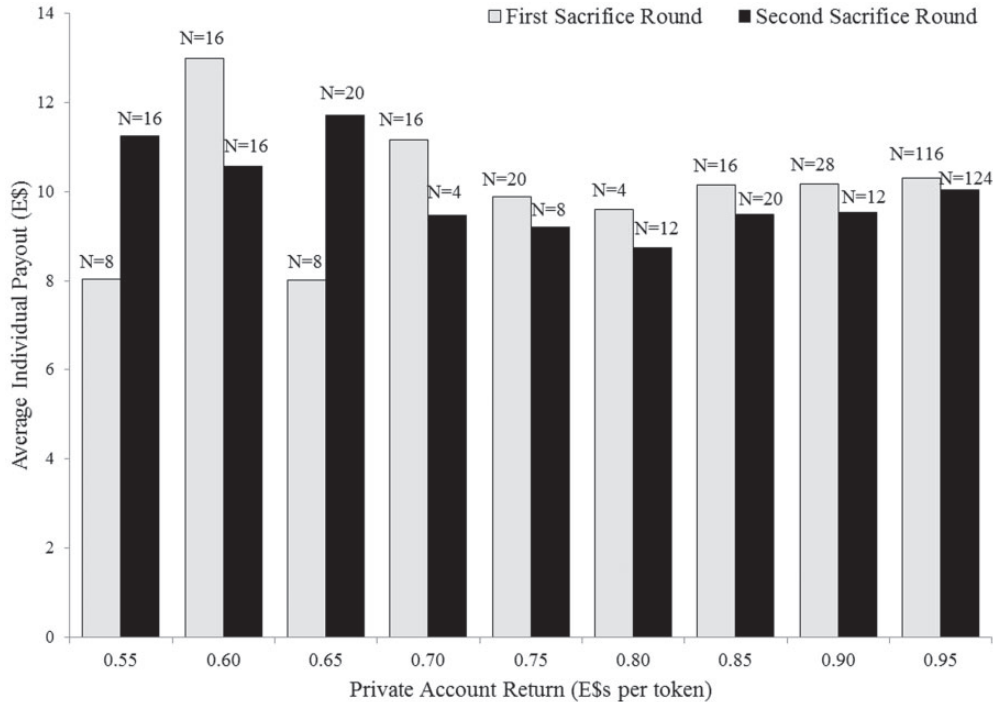


FIGURE 6
Average individual payout

the spectrum of sacrifice. The negative trend is small but statistically significant in both first and second sacrifice rounds.³⁵

4.3. Screening versus substitution

The prohibitions and mandates that discourage free-riders also encourage higher contributions of those who remain by lowering their productivity outside the group. In fact, this substitution effect can sometimes improve welfare in and of itself—thereby making sacrifice a viable strategy even for pooled or homogeneous populations that have no use for screening (Iannaccone, 1992).

Standard VCM experiments confirm that subjects do substitute, responding as expected to changing relative rates of return to private versus group investment (Ledyard, 1995: 149–150). To separate pure substitution effects from those of screening, we ran three “control” sessions. As in the Inexperienced sessions, subjects first played a Sacrifice round, then a Normal round, and finally a second Sacrifice round. But in contrast to all other sessions, these “control” sessions gave subjects no opportunity to choose their levels of sacrifice. Instead, we randomly assigned

H-types (if this effect is strong enough). This is similar to the model in Levy and Razin (2010), where agents are better off when they can coordinate with others on religious actions. Since this abstracts from our main point and is not tested in the lab, we do not push this point any further.

35. The Cuzick trend test yields p -values less than 0.05 and 0.005 for the first and second Sacrifice rounds. Dividing the first and second rounds results by Experienced versus Inexperienced subjects yields a statistically significant decreasing trend only with Inexperienced subjects (with $p < 0.08$ in the first round and $p = 0.001$ in the second, but $p > 0.4$ for Experienced subjects in both rounds).

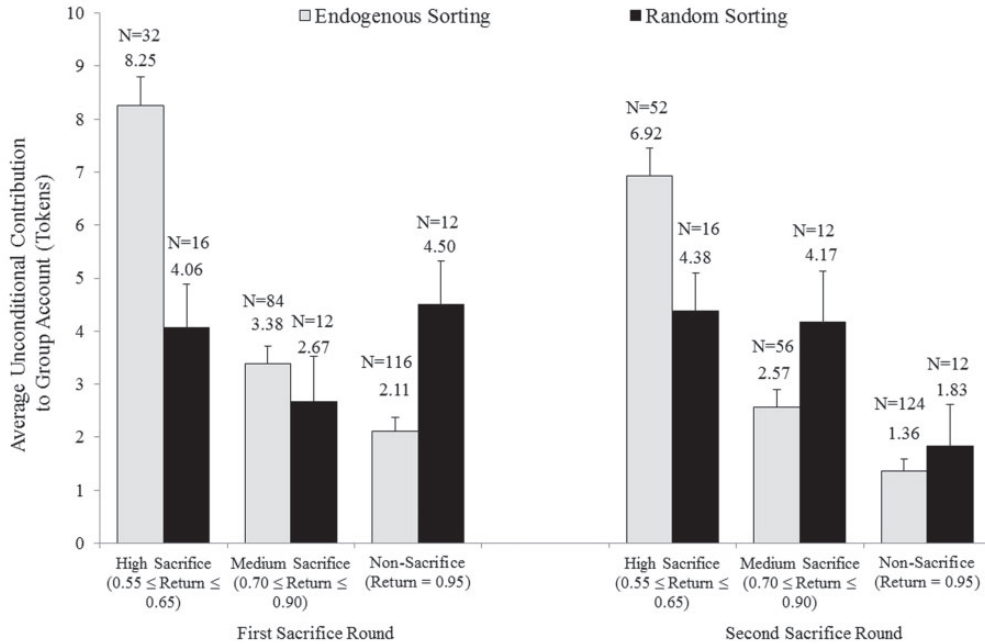


FIGURE 7

Endogenous versus random comparison, average unconditional contribution, multiple groups

them to “High,” “Medium,” or “Low” sacrifice groups, which paid private return rates of 0.60, 0.75, or 0.95.³⁶

Figure 7 compares average contributions in randomly sorted “control” sessions versus endogenously sorted regular sessions. Although price effects alone may have some effect, they seem relatively weak. In control sessions, overall rates of giving in both the first and second sacrifice round combined decline as expected from \$4.2 in “High” sacrifice groups, to \$3.4 in “Medium,” to \$3.2 in “Low,” but no clear pattern stands out when rounds are analysed separately.³⁷

In contrast, the relationship between rates of giving and rates of sacrifice is much stronger and much more consistent for endogenously sorted subjects. The pattern becomes stronger still when, as in Figure 2, we graph each level of sacrifice separately. Consistent with Prediction 5, endogenous sorting yields larger average unconditional contributions than random sorting in high-sacrifice groups (Mann–Whitney two-tailed, $p < 0.01$ in the first and second sacrifice rounds). Conversely, endogenous sorting leads to lower unconditional contributions in non-sacrifice groups ($p < 0.01$ in the first sacrifice round but $p = 0.16$ in the second).³⁸

36. We chose the “Inexperienced” ordering in control sessions so we would see decisions with exogenously set private return rates, randomly matched group members, and no prior experience with the game—thereby controlling for both screening and learning effects. Figure 7 compares control subjects to all other subjects, but the same pattern emerges when limited only to Inexperienced subjects.

37. The lack of trend in separated rounds may be due in part to the small number of control session subjects, but even if we take the pooled results as definitive, the magnitude of the giving-versus-sacrifice trend is much smaller with random sorting than with endogenous sorting.

38. The endogenous–exogenous difference is statistically insignificant ($p > 0.10$) in both the first and second sacrifice rounds in medium sacrifice groups.

So while it remains likely that sacrifice induces some degree of substitution in the absence of screening (particularly in light of evidence from standard VCM experiments, not to mention the law of demand), it seems clear that screening effects greatly strengthen the association between sacrifice and club production—in our experiment and probably also in the world.

5. CONCLUSION

Sacrifice has long been thought to foster commitment in religions, communes, fraternities, clans, gangs, military units, subcultures, social movements, political organizations, and countless other groups. More recently, economists have argued that sacrifice often works by raising the cost of alternative activities, so as to screen out people most likely to free-ride and strengthen involvement among those who remain. The economic literature on seemingly inefficient prohibitions and mandates now includes formal theory, survey research, case studies, and simulations covering a wide range of issues and applications. Yet theory and field work cannot readily separate screening and substitution from deviant tastes, unusual beliefs, effective socialization, rigid tradition, cognitive dissonance, and outright irrationality.

We have tried to carry the economic theory of sacrifice into the lab, implementing it in a way that leaves little room for other effects, most especially the religious activities and institutions that feature prominently in past research. To keep our experiment as clean as possible, we worked with linear VCMs, the simplest and most frequently studied club goods game. We altered the standard setup to let players select themselves into groups with reduced rates of return to private investment. We also used a popular method for typing free-riders, conditional cooperators, and others.

The results are striking. When given the option to effectively destroy a portion of their private productivity, with no compensation but the prospect of being grouped with like-minded others, more than half of all cooperators did so. But free-riders overwhelmingly rejected the offer, thus enabling cooperative sacrificers to greatly increase their joint club output. Nor was this the only effect; sacrifice and contributions correlated across the entire spectrum of sacrifice. Contributions averaged around 85% of total endowment in the most costly groups, 35% in the moderately costly groups, and 15% in the least costly groups. The screening of free-riders and substitution towards club production was so great that sacrifice groups earned as much overall as no-sacrifice groups (supplemented, we presume, by the glow of good citizenship). Finally, screening was not so strong as to suggest mere type-based habits. Nearly one in five free-riders did choose sacrifice groups, presumably hoping that the kindness of strangers would more than offset the guaranteed loss on their own (privately directed) resources.

We of course welcome both replications and reanalysis of our data. But we especially encourage extensions. An easy first step would be sacrifice VCMs with many repetitions and intermittent options to change groups—in part to determine whether the mechanism can limit contribution decay or progressive infiltration. Other extensions might test additional features of the sacrifice model, including its predictions concerning group dynamics and demographics (Iannaccone, 1992, 1994; Makowsky, 2012). For obvious reasons, high-sacrifice groups more often attract people with limited outside opportunities—including the poor, recent immigrants, and the young. This tendency could be tested by giving subjects different base rates of private productivity. Sacrifice mechanisms could also be introduced to other standard experiments, including trust and ultimatum games.

We have contrasted sacrifice with more common solutions to collective action problems, including exclusion, fees, rewards, punishment, and expressions of approval. But real-world clubs routinely embrace multiple methods, and more realistic experiments might do the same. A combination of sacrifice and other approaches might prove especially effective in experiments

where information about the other members' inputs can only be observed intermittently, noisily, or by random others.

Sacrifice is no panacea. It is, in fact, never more than a second-best solution to collective action problems. But its strength lies precisely in its capacity to continue working where other methods fail.

Acknowledgments. We are grateful for comments received in seminars at Cal State Fullerton, Chapman University, George Mason University, the 2009 ASREC meetings, the 2010 meetings of the Public Choice Society, and the 2010 meetings of the Western Economic Association. We thank ICES for use of its laboratory facilities, and the John Templeton Foundation, IFREE, the Mercatus Center, and Chapman University for funding support. Aimone was funded in part by Brooks King-Casas and Makowsky received funding from Joshua Epstein's NIH Director's Pioneer Award, Number DP1OD003874 from the Office of the Director, National Institutes of Health.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- ABRAMITZKY, R. (2008), "The Limits of Equality: Insights from the Israeli Kibbutz", *Quarterly Journal of Economics*, **123**, 1111–1159.
- ABRAMITZKY, R. (2009), "The Effect of Redistribution on Migration: Evidence from the Israeli Kibbutz", *Journal of Public Economics*, **93**, 498–511.
- AHN, T. K., MARK ISAAC, R. and SALMON, T. C. (2008), "Endogenous Group Formation", *Journal of Public Economic Theory*, **10**, 171–194.
- ANDERSON, C. M. and PUTTERMAN, L. (2006), "Do Non-Strategic Sanctions Obey the Law of Demand? The Demand for Punishment in the Voluntary Contribution Mechanism", *Games and Economic Behavior*, **54**, 1–24.
- ANDREONI, J. (1990), "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving", *Economic Journal*, **100**, 464–477.
- ARONSON, E. and MILLS, J. (1959), "The Effect of Severity of Initiation on Liking for a Group", *Journal of Abnormal and Social Psychology*, **59**, 177–181.
- BERMAN, E. (2000), "Sect, Subsidy, and Sacrifice: An Economist's View of Ultra-Orthodox Jews", *Quarterly Journal of Economics*, **65**, 905–953.
- BERMAN, E. (2009). *Radical, Religious, and Violent: The New Economics of Terrorism* (Cambridge: MIT Press).
- BERNDT, C. E. H. (2007), "High-Tension Religion and the Jewish Peddler", *American Journal of Economics and Sociology*, **66**, 1005–1027.
- BOHNET, I. and KUBLER, D. (2005), "Compensating the Cooperators: Is Sorting in the Prisoner's Dilemma Possible?", *Journal of Economic Behavior and Organization*, **56**, 61–76.
- BOSE, F. (2010), "The Green Marketplace: Applying a Model of Church and Sect to the Environmental Movement", Meetings of the Public Choice Society. Monterey, CA.
- BRAM, U. (2010), "The Big Screen: Club Theory and the Chinese Communist Party" (SSRN Working Paper No. 1550189).
- BRANDTS, J. and SCHRAM, A. (2008), "Cooperation in VCM Experiments: Results Using the Contribution Function Approach", *Handbook of Experimental Economics Results*, **1**, 825–830.
- BUCHANAN, J. M. (1965), "An Economic Theory of Clubs", *Economica*, **32**, 1–14.
- BURLANDO, R. M. and GUALA, F. (2005), "Heterogeneous Agents in Public Goods Experiments", *Experimental Economics*, **8**, 35–54.
- CARVALHO, J.-P. (2013), "Veiling", *Quarterly Journal of Economics*, **128**, 337–370.
- CHAN, K. S., MESTELMAN, S. and MULLER, R. A. (2008), "Voluntary Provision of Public Goods", *Handbook of Experimental Economics Results*, **1**, 831–835.
- CHARNESS, G. and YANG, C.-L. (2010), "Public Goods Provision with Voting for Exclusion, Exit, and Mergers: An Experiment" (MIMEO).
- CHATTOE, E. (2006), "Using Simulation to Develop and Test Functionalist Explanations: A Case Study of Dynamic Church Membership", *British Journal of Sociology*, **57**, 379–397.
- CHAUDHURI, A. (2011), "Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature", *Experimental Economics*, **14**, 47–83.
- CHEN, D. (2010), "Club Goods and Group Identity: Evidence from Islamic Resurgence during the Indonesian Financial Crisis", *Journal of Political Economy*, **118**, 300–352.
- CIMINO, A. (2011), "The Evolution of Hazing: Motivational Mechanisms and the Abuse of Newcomers", *Journal of Cognition and Culture*, **11**, 241–267.
- CORICELLI, G., FEHR, D. and FELLNER, G. (2004), "Partner Selection in Public Goods Experiments", *Journal of Conflict Resolution*, **48**, 356–378.

- CORNES, R. and SANDLER, T. (1984), "Easy Riders, Joint Production, and Public Goods", *Economic Journal*, **94**, 580–598.
- CRUMPLER, H. and GROSSMAN, P. J. (2008), "An Experimental Test of Warm Glow Giving", *Journal of Public Economics*, **92**, 1011–1021.
- D'AMICO, D. J. (2008), "Tattoo Prohibition Behind Bars: The Case for Repeal", *Journal of Private Enterprise*, **23**, 113–134.
- FESTINGER, L. and CARLSMITH, J. M. (1959), "Cognitive Consequences of Forced Compliance", *Journal of Abnormal and Social Psychology*, **56**, 276–278.
- FINKE, R. and STARK, R. (1992), *The Churching of America, 1776–1990: Winners and Losers in our Religious Economy* (New Brunswick: Rutgers University Press).
- FISCHBACHER, U., GÄCHTER, S. and FEHR, E. (2001), "Are People Conditionally Cooperative? Evidence from a Public Goods Experiment", *Economics Letters*, **71**, 397–404.
- FISCHBACHER, U., GÄCHTER, S. and QUERCIA, S. (2012), "The Behavioral Validity of the Strategy Method in Public Good Experiments", *Journal of Economic Psychology*, **33**, 897–913.
- GÄCHTER, S. and THÖNI, C. (2005), "Social Learning and Voluntary Cooperation Among Like-Minded People", *Journal of the European Economic Association*, **3**, 303–314.
- GRIMM, V. and MENGEL, F. (2009), "Cooperation in Viscous Populations – Experimental Evidence", *Games and Economic Behavior*, **66**, 202–220.
- GUNNTHORSDDOTTIR, A., HOUSER, D. and MCCABE, K. (2007), "Disposition, History and Contributions in Public Goods Experiments", *Journal of Economic Behavior & Organization*, **62**, 304–315.
- GUNNTHORSDDOTTIR, A., VRAGOV, R., SEIFERT, S. and MCCABE, K. A. (2010a), "Near-Efficient Equilibria in Collaborative Meritocracies", *Journal of Public Economics*, **94**, 987–994.
- GUNNTHORSDDOTTIR, A., VRAGOV, R. and SHEN, J. (2010b), "Tacit Coordination in Contribution-Based Grouping with Two Endowment Levels", in Isaac, R. M. and Norton, D. A. (eds) *Charity with Choice: Research in Experimental Economics*, Vol. 13 (Bingley, U.K.: Emerald Group Publishing Limited) 13–75.
- HANSON, G. H. and XIANG, C. (2011), "Exporting Christianity: Governance and Doctrine in the Globalization of US Denominations" (NBER Working Paper 16964).
- HOUSER, D., XIAO, E., MCCABE, K. and SMITH, V. (2008), "When Punishment Fails: Research on Sanctions, Intentions and Non-Cooperation", *Games and Economic Behavior*, **62**, 509–532.
- IANNACCONI, L. R. (1992), "Sacrifice and Stigma: Reducing Free-Riding in Cults, Communes, and Other Collectives", *Journal of Political Economy*, **100**, 271–291.
- IANNACCONI, L. R. (1994), "Why Strict Churches are Strong", *American Journal of Sociology*, **99**, 1180–1211.
- IANNACCONI, L. R. (2006), "The Market for Martyrs", *Interdisciplinary Journal of Research on Religion*, **2**.
- IANNACCONI, L. R. and BERMAN, E. (2006), "Religious Extremism: The Good, the Bad, and the Deadly", *Public Choice*, **128**, 109–129.
- IRIBERRI, N. and REY-BIEL, P. (2009), "Elicited Beliefs and Social Information in Modified Dictator Games: What Do Dictators Believe Other Dictators Do?" (SSRN Working Paper No. 1374287).
- JAMES, R. N. and SHARPE, D. L. (2007), "The 'Sect Effect' in Charitable Giving: Distinctive Realities of Exclusively Religious Charitable Givers", *American Journal of Economics and Sociology*, **66**, 697–726.
- KOYAMA, M. and CARVALHO, J.-P. (2011), "Development and Religious Polarization: The Emergence of Reform and Ultra-Orthodox Judaism" (George Mason University Working Paper in Economics, No. 11-45).
- LEDYARD, J. O. (1995), "Public Goods: A Survey of Experimental Research", in Kagel, J. and Roth, A. E. (eds) *Handbook of Experimental Economics* (Princeton: Princeton University Press).
- LEMBKE, A. (2012), "Sacrifice and Stigma in Alcoholics Anonymous: Applying Iannaccone's Model to Self-Help Organizations for Addiction" (MIMEO, Stanford).
- LEVY, G. and RAZIN, R. (2010), "Religious Organizations" (MIMEO).
- MAKOWSKY, M. D. (2011), "A Theory of Liberal Churches", *Mathematical Social Sciences*, **61**, 41–51.
- MAKOWSKY, M. D. (2012), "Emergent Extremism in a Multi-Agent Model of Religious Clubs", *Economic Inquiry*, **50**, 327–347.
- MASCLET, D., NOUSSAIR, C., TUCKER, S. and VILLEVAL, M.-C. (2003), "Monetary and Nonmonetary Punishment in the Voluntary Contributions Mechanism", *American Economic Review*, **93**, 366–380.
- MCBRIDE, M. (2007), "Why Churches Need Free-Riders: Religious Capital Formation and Religious Group Survival" (Working Paper #060722, University of California-Irvine, Department of Economics).
- MONTGOMERY, J. (1996), "Dynamics of the Religious Economy: Exit, Voice and Denominational Secularization", *Rationality and Society*, **8**, 81–110.
- NOUSSAIR, C. and TUCKER, S. (2005), "Combining Monetary and Social Sanctions to Promote Cooperation", *Economic Inquiry*, **43**, 649–660.
- OLSON, M. (1965), *The Logic of Collective Action: Public Goods and the Theory of Groups* (Cambridge: Harvard University Press).
- ORBELL, J. M. and DAWES, R. M. (1993), "Social Welfare, Cooperators' Advantage, and the Option of Not Playing the Game", *American Sociological Review*, **58**, 787–800.
- OSTROM, E. (1990), *Governing the Commons: The Evolution of Institutions for Collective Action* (Cambridge: Cambridge University Press).
- PAGE, T., PUTTERMAN, L. and UNEL, B. (2005), "Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency", *Economic Journal*, **115**, 1032–1053.

- POSNER, E. A. (2002), *Law and Social Norms* (Cambridge: Harvard University Press).
- RAPOPORT, H., and WEISS, A. (2002). "In-Group Cooperation in a Hostile Environment: An Economic Perspective on Some Aspects of Jewish Life in (Pre-Modern) Diaspora" (IZA Discussion Paper No. 483).
- RICHARDSON, G. and MCBRIDE, M. (2009), "Religion, Longevity, and Cooperation: The Case of the Craft Guild", *Journal of Economic Behavior & Organization*, **71**, 172–186.
- SHIMIZU, H. (2011), "Social Cohesion and Self-Sacrificing Behavior", *Public Choice*, **149**, 427–440.
- SOSIS, R. (2000), "Religion and Intragroup Cooperation: Preliminary Results of a Comparative Analysis of Utopian Communities", *Cross-Cultural Research*, **34**, 70–87.
- SOSIS, R., KRESS, H. C. and BOSTER, J. S. (2007), "Scars for War: Evaluating Alternative Signaling Explanations for Cross-Cultural Variance in Ritual Costs", *Evolution and Human Behavior*, **28**, 234–247.
- STARK, R. (1996), *The Rise of Christianity: A Sociologist Reconsiders History* (Princeton: Princeton University Press).
- STARK, R. and IANNACCONE, L. R. (1997), "Why Jehovah's Witnesses Grow so Rapidly", *Journal of Contemporary Religion*, **12**, 133–157.
- SWOPE, K. J. (2002), "An Experimental Investigation of Excludable Public Goods", *Experimental Economics*, **5**, 209–222.
- THIES, C. F. (2000), "The Success of American Communes", *Southern Economic Journal*, **67**, 186–199.
- XIAO, E. and HOUSER, D. (2005), "Emotion Expression in Human Punishment Behavior", *Proceedings of the National Academy of Science*, **102**, 7398–7401.