

http://researchcommons.waikato.ac.nz/

Research Commons at the University of Waikato

Copyright Statement:

The digital copy of this thesis is protected by the Copyright Act 1994 (New Zealand).

The thesis may be consulted by you, provided you comply with the provisions of the Act and the following conditions of use:

- Any use you make of these documents or images must be for research or private study purposes only, and you may not make them available to any other person.
- Authors control the copyright of their thesis. You will recognise the author's right to be identified as the author of the thesis, and due acknowledgement will be made to the author where appropriate.
- You will obtain the author's permission before publishing any material from the thesis.

Cooperation, incentives and punishment in common pool resource management

A thesis
submitted in partial fulfilment
of the requirements for the degree
of
Master of Management and Sustainability
at
The University of Waikato
by
OLIVER J. PARSONS



THE UNIVERSITY OF
WAIKATO
Te Whare Wānanga o Waikato

2016

Abstract

Managing water quality is of critical interest to policy-makers in New Zealand and globally. In particular, the management of diffuse nitrate losses from agriculture is a major policy challenge that remains largely unsolved. In a policy environment with a strong emphasis on collaboration, consideration has been given to whether groups of resource users can address the issue collectively at a catchment scale. Although water quality is not a traditional common pool resource problem, in the presence of some form of external threat or regulation associated with degradation, the diffuse loss of nitrate to waterways is transformed into a common pool resource problem. This social dilemma is characterised by strong incentives to defect and challenging levels of complexity. However, these challenges are not insurmountable.

This study uses economic experiments to present groups of participants with a stylised version of the problem faced by resource users in the field. Previous experiments in this area have found sanctions to be of limited effectiveness in this environment in the absence of communication. We extend previous studies by adjusting the ratio of the cost of sanctions to their effectiveness, and by varying the overall payoff environment. We find that more effective sanctions are able to improve environmental outcomes and stabilise the level of resource use over time, even in the absence of communication. However, there are significant costs associated with the use of sanctions which mean that overall welfare is not enhanced in the absence of communication.

Due to the complexity of common pool resource management, communication and agreement about norms for resource use levels is important to support groups in optimising environmental outcomes and profitability. Communication is a necessary condition for maximising welfare when navigating complexity. However, communication is not necessary to improve environmental outcomes or sustain a level of resource use over time.

While some self-selected groups of highly-cooperative individuals may be able to sustain cooperation in the absence of sanctions, these arrangements are unlikely to be stable over time. The presence of defectors means that some form of recourse to punishment is necessary to support the continued sustainable management of a resource. The strong incentives for defection in common pool resource management, including collective management of water quality, can be overcome by sufficiently strong sanctions.

Our experimental findings suggest that sanctions administered by peers can be effective in improving environmental outcomes and stabilising the level of resource use over time, preventing the decline in cooperation that is evident where individuals are able to profit from defection. The cost and effectiveness of sanctions has a strong interaction with the nature of the incentive environment. In designing institutions for the collective management of nitrate loads to water bodies, regulators and groups of resource users must consider the incentives for defection that are driven by the degree of coupling of nitrate losses and profitability in farm systems. In order to be successful, group members must have access to low-cost sanctions that are designed with the incentive environment in mind. If institutions can be designed in such a way as to ensure that cheaters do not prosper from defection, the prospects for collective management of water quality are positive.

Acknowledgments

Profound thanks are due to a number of persons who have supported me in the production of this thesis. First, I would like to thank DairyNZ for providing the funding which made this research possible. Second, my supervisors Steven Tucker and Graeme Doole for their continual challenging of my assumptions and conscientious push to improve the quality of the research. Third, my supportive and understanding friends, parents, siblings and wider whānau. Last but not least, special thanks are due to the incomparable Rosie Batt.

Contents

Abstract iii

Acknowledgments v

Contents vii

List of tables xi

List of figures xiii

1. Introduction and background 1

1.1 Freshwater in New Zealand 2

1.2 Agriculture and non-point source pollution 3

1.3 Collective approaches to managing freshwater quality 6

1.4 Motivation 8

1.5 Aims 10

2. Literature Review 11

2.1 Non-point source pollution and common pool resources 11

2.2 Experimental economics and environmental policy 15

2.3 Social dilemma game environments 19

Linear Voluntary Contribution Mechanism (VCM) 20

Non-Linear VCM 21

Common Pool Resource (CPR) Game 21

Pollution games and ambient schemes 23

2.4 Incentives – group size, marginal per-capita return and rivalry 24

2.5 Sanctions 26

2.7 The cost of sanctions 33

2.8 Norms and complexity 36

2.9 Covenants and communication 40

2.10 Critical research gaps of interest 42

3. Method 45

3.1 Hypothesis 45
3.1 Location and participants 46
3.2 Economic setting 47
3.4 Punishment 50
Explanation 1: Punishment levels respond to damage levels 52
Explanation 2: Prospering cheaters 53
3.5 Treatments 54
3.6 Game Theoretic Predictions 55
Low payoff environment 55
High payoff environment 57
4. Results 59
4.1 Extraction levels 59
Result one: Stronger punishment results in lower extraction levels 59
Result two: Stronger punishment inhibits decline in cooperation over time . 62
4.3 Punishment – use 63
Result three: More effective punishment and higher incomes generate greater demand for punishment (punishment is a normal good) 63
Result four: More effective punishment is used to increasingly target defectors 66
4.4 Punishment – impact 67
Result five: The direct impacts of strong punishment interact with the indirect effects 67
Result six: More effective punishment can remove the capability for defectors to earn more than co-operators 69
4.2 Efficiency 71
Result seven: The cost of increasingly severe punishment offsets the welfare gains of lower extraction (at least in the short-term) 72
5. Discussion 77

5.1 Levels of cooperation and extraction 77
Hedging 79
Confusion 80
5.2 Incentives 81
5.3 Punishment 83
The cost of punishment 83
Income effects 84
5.4 Relationship between incentives, punishment ratio and cooperation 85
5.5 Efficiency 88
5.6 Implications for collective management of water quality 89
6. Summary and Conclusion 93
6.1 The relationship between incentives, punishment and cooperation in common pool resource games 93
6.2 Limitations and future extensions 94
6.3 Concluding remarks 95
7. References 97
8. Appendix: Instructions for a 1:9 punishment experiment 111

List of tables

Table 1: Excludability and rivalry for different types of property 13
Table 2: Property rights regimes in relation to resource management 14
Table 3: Outline of key features of the experimental environments 49
Table 4: Differences in incentives between linear VCM and CPR environment and associated punishment ratios 54
Table 5: Treatments 55
Table 6: Nash equilibrium predictions for low payoff games 57
Table 7: Social optimum predictions for low payoff games 57
Table 8: Nash equilibrium predictions for high payoff games 58
Table 9: Social optimum predictions for high payoff games 58
Table 10: Mean extraction levels 60
Table 11: Differences in extraction between treatments 60
Table 12: Correlations in extraction over time 62
Table 13: Mean payoff reductions received 64
Table 14: Differences in punishment use 66
Table 15: Correlations between punishment and change in extraction in the following period 69
Table 16: Differences in earnings due to punishment 71
Table 17: Mean efficiency 74
Table 18: Differences in efficiency 75

List of figures

Figure 1: Relative suitability of different regulatory approaches 5
Figure 2: The group benefit functions used in the CPR experiments. The maximum of each curve represents the maximum yield from resource extraction. 48
Figure 3: Comparison of payoffs from defection in linear VCM and low payoff CPR games 51
Figure 4: Comparison of ratios of gain to damage caused by defection in linear VCM and low payoff CPR environments 52
Figure 5: Mean extraction - low payoff treatments 61
Figure 6: Mean extraction comparing high and low payoff environments 61
Figure 7: Mean contributions - low payoff treatments 63
Figure 8: Payoff reductions received in each period 65
Figure 9: Demand for punishment at 1:3, 1:5 and 1:9 effectiveness ratios 65
Figure 10: Payoff reductions received 67
Figure 11: Payoff reductions received and mean reaction 68
Figure 12: Earnings from defection - all low payoff treatments 70
Figure 13: Earnings from defection - comparison with high payoff treatment 70
Figure 14: Mean efficiency relative to social optimum 73
Figure 15: Influence diagram for impacts of punishment effectiveness 88

1. Introduction and background

Freshwater management is a critical issue for regional and national governments worldwide. Increasing pressure on the resource from extraction and pollution has caused management to reach a crisis point in many areas. As freshwater is of critical economic, social and cultural importance to human societies, resolving approaches to management is important and urgent. In New Zealand, the *National Policy Statement for Freshwater Management* was introduced in 2011, requiring limits to be set for both quality and quantity for all freshwater bodies. This was strengthened in 2014 with the addition of a set of national bottom lines, setting minimum standards for water quality (New Zealand Government, 2014). Regional councils are now required to set limits and to manage activities to achieve these limits. However, how councils are expected to achieve this is not prescribed. Indeed, the methods being introduced to achieve limits are as diverse as the freshwater management issues that present themselves throughout the country (Ministry for the Environment, 2015a).

Embedded within this policy setting is a shift towards more collaborative approaches to setting policy and rules, as well as more collaborative approaches to management (Ministry for the Environment, 2015b). Given the strong impetus towards group problem solving, it is perhaps unsurprising that collective approaches towards freshwater management have gained considerable attention. However, policy-makers and resource users are now utilising models of collective management in areas where they have been previously untested, such as in the broad area of water quality. Although the commitment towards collaboration is generally strong in terms of both politics and implementation, the approach is still very new. Much remains unknown about how this new collaborative approach to environmental governance will be sustained.

1.1 Freshwater in New Zealand

Freshwater quality has many different components. Clarity, suspended sediment, nutrient levels, algal growth, dissolved oxygen, biodiversity and the impact of endangered species are all affected by different drivers that interact dynamically (Ballantine & Davies-Colley, 2013). In the relatively short history of New Zealand, the dominant drivers and attendant states have changed dramatically. Vegetation clearance by Māori is likely to have had some impact by accelerating erosion, but this would have been relatively modest in terms of ecological disturbance. Significant degradation began with European settlement. Settlements discharged untreated sewage directly to surface waters (and this still occurs in many areas as significant rain tends to overwhelm existing infrastructure). Mining booms resulted in toxic contamination with persistent long-term effects for some areas. Also, across the entire country, deforestation caused accelerated erosion, which has supported a legacy of sedimentation and flooding. Government policy focused on settlement and economic development accelerated and worsened these impacts. For example, subsidies for clearing erosion-prone hill country for unprofitable farming operations persisted well into the 1980s (Parliamentary Comissioner for the Environment, 2012).

The dominant change to freshwater environments over the last half-century has been increasing levels of nutrients (Ballantine & Davies-Colley, 2013). Rapid expansion of towns and factories has resulted in increasing numbers of point-source discharges to freshwater bodies. The use of phosphate fertiliser since the 1960s enabled an enormous increase in agricultural productivity, increasing the carrying capacity of pastures several times over. However, agricultural land use intensity has increased steadily ever since, supporting ever-greater levels of production, but also waste in the form of diffuse nutrient losses (Parfitt et al., 2012). Since the introduction of the *Resource Management Act 1991*, significant progress has been made in reducing the impacts of many point-source discharges on freshwater. While farming practices have also improved in terms of improving nutrient use efficiency, their relative contribution to the nutrient enrichment of water bodies, compared to point sources, has grown. In addition, the trend of

increasing intensity has driven increases in nitrogen levels in many areas (Parfitt et al., 2012; Parliamentary Comissioner for the Environment, 2012).

1.2 Agriculture and non-point source pollution

Nutrients (particularly nitrogen and phosphorus) are essential to supporting freshwater ecosystems. However, excessive levels can cause undesirable algal blooms, which increase the diurnal oscillation in dissolved oxygen levels, making the water unsuitable for fish (Matheson, Quinn, & Hickey, 2012). At very high levels, nitrates can also have direct toxic effects on freshwater fauna and make it unsuitable for human consumption (Hickey & Martin, 2009; World Health Organization, 2004). Following early success with the regulation of point-source discharges, attention has now turned to the management of diffuse losses of nitrogen and phosphorus from agriculture, often referred to in the literature as non-point source pollution (Shortle & Horan, 2001).

The most effective management approaches for nitrogen and phosphorus losses from agriculture are very different. As phosphorus binds to soil particles, there are a range of options at a farmer's disposal which enable effective mitigation (for example, through stock exclusion and riparian planting) that are largely decoupled from production. Nitrate-nitrogen, however, is much more costly to manage (Monaghan et al., 2007). While higher nitrogen use efficiency allows for some increase in production within given environmental boundaries, rates of nitrate leaching will always bear some relationship to the overall intensity of production in a catchment (Clapcott, Young, Goodwin, Leathwick, & Kelly, 2011).

Because of this coupling of nitrogen loss and production, the mitigation of nitrogen involves a more conventional convex relationship between abatement and abatement cost, compared to phosphorus. This makes management much more challenging, as it is not a simple matter of prescribing a particular technology to fix the problem. Where any significant level of mitigation is required beyond reasonable efficiency gains, there will be significant implications for profitability (Doole & Kingwell, 2015). Given the incentive structure

surrounding the abatement of nitrogen loss from farms, some form of on-going management is likely to be required, whether it is through regulation of individuals or groups.

Expense is not the only challenge. There are several characteristics of agricultural non-point source pollution that make it difficult to regulate individual firms. In particular, the existence of multiple polluters, information asymmetries, complex transmission pathways and stochastic environmental influences make non-point source pollution problems fraught with uncertainty and unobservability (Shortle & Horan, 2001). Agricultural non-point source pollution is also characterised by natural variability, and is often relatively local scale (Romstad, 2003; Xepapadeas, 1992). The impact of natural variability is a product of processes such as weather and climate, as well as the responses of firms to these changes in conditions. Measurement of firm-level emissions under these conditions becomes prohibitively expensive (Suter, Vossler, Poe, & Segerson, 2008). This has hampered efforts to address the problem for decades (Organisation for Economic Cooperation and Development (OECD), 2012).

The instruments typically used to address pollution problems can be characterised into three broad classes:

- Performance-based (focused on actual outputs such as kilograms of nitrogen leached),
- Management-based (focused on actions taken such as irrigating effluent from dairy sheds in a way that avoids runoff), and
- Technology-based (implementing particular technologies such as riparian-buffer planting).

Coglianese and Lazer (2003) propose a useful framework for showing the different strengths of these three approaches to environmental management, represented pictorially in figure 1. Where the polluting firms are highly homogenous with respect to a particular problem, it can be effective to target regulation towards a particular technology. Where firms are more diverse, regulation tends to focus on management or performance-based instruments.

Because of the information problems associated with non-point source pollution, regulation has typically focused on the implementation of technologies or management practices that are correlated with environmental impacts – for example, fertilisation, manure spreading or tillage (Doole, 2010; OECD, 2012).

Figure 1: Relative suitability of different regulatory approaches

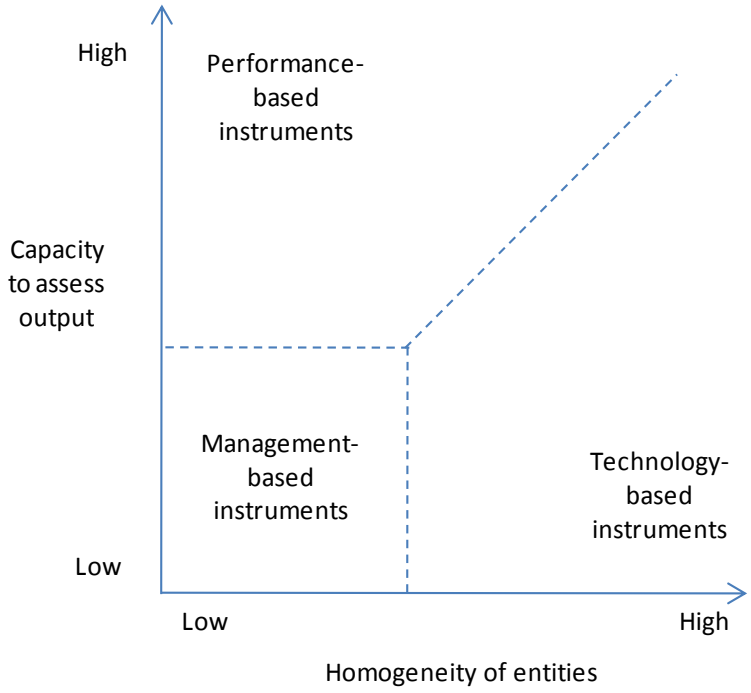

(Coglianese & Lazer, 2003)

Technology and management-based regulatory approaches have been effective in some areas, but in others, policies have failed (for instance, in controlling nitrate leaching) (OECD, 2007). Although inputs and pollutant loads are often correlated, these technologies and practices are not sufficiently connected to the level of observed nitrate loss to be an effective means of control. For example, fertiliser use may explain only 30 per cent of nitrate losses from a given farm system, with large variation attributable to the efficiency of input use or natural variability. More direct, performance-based approaches would be preferable from both an environmental and an efficiency perspective (Doole, 2012; Romstad, 2003). Moreover, the regulation of inputs has to be very broad-ranging if

producers are not to adapt through substituting regulated factors of production with other polluting inputs. For example, banning nitrogen fertiliser can reduce nitrogen loss from farming systems, but can also provide a strong incentive to increase supplement feeding, which erodes any environmental benefit accruing to the ban placed on nitrogen application (Doole, 2010).

By taking a performance-based approach to regulation, management behaviour also becomes directly connected to the underlying incentive structure that is driven by the abatement-cost relationship. Where these costs are large, the incentives to not comply with regulation become significant. In the context of group management that is being considered in some areas, the incentives to defect, to the detriment of the welfare of the resource management group as a whole, become significant.

1.3 Collective approaches to managing freshwater quality

Numerous examples exist of using a collective approach to the management of water quantity for irrigation (Cox & Ross, 2011; Meinzen-Dick, 2007). Indeed, many small irrigation schemes in New Zealand use a collective approach. However, such approaches are not uniformly successful and there are a number of limitations (Meinzen-Dick, 2007). Water quantity, in contrast to water quality, is relatively easily defined as a resource system. Consequently, this has made the establishment of user boundaries somewhat easier in that the supply can always be turned off. The highly visible act and results of irrigation also make monitoring relatively straightforward (even more so with the modernisation of irrigation systems and the incorporation of water metering and telemetry). However, increased difficulty does not mean it is impossible to apply a common property approach to water quality.

A number of collective efforts by farmers to address water quality issues have already occurred in New Zealand. The *Aorere Catchment Project* and *Lake Rerewhakaaitu Catchment Action Plan* are examples (Bay of Plenty Regional Council, 2012; New Zealand Landcare Trust, 2012). In the case of the Aorere

project, a number of dairy farmers made collective efforts to reduce bacterial impacts from farm dairy effluent on local mussel farms. The Rerewhakaaitu project included measures in relation to effluent, riparian planting and nutrient management. However, no such group currently exists focused chiefly on the management of nitrate. While the large incentives for defection that characterise the problem undoubtedly make collective management of nitrate losses more challenging, there are a number of reasons to suppose that a group approach might be desirable.

A particular feature of the non-point source pollution problem is that it often occurs at a relatively local scale, where resource users tend to be in relatively close contact. This has two important consequences. First, the scale and potential for group cohesiveness lend the situation towards collective action or group management (Romstad, 2003). Second, users within the group may have an information advantage relative to regulators, due to their ability to partially observe their neighbours' behaviour (for example, noticing effluent spills, poor cultivation practices or additional, undeclared bought-in feeds and nutrients) (Seabright, 1993). In managing an environmental problem that is so fundamentally characterised by information problems, this information advantage could mean group management approaches are more effective than controls on individual firms by the state. Finally, in the case of mitigating nitrogen loss, there is considerable scope to change the abatement curves do be less costly through increasing farmer skill (Doole & Kingwell, 2015). However, achieving this profitable abatement through better nutrient management can be difficult, while the profit gains may be marginal (Pannell, 2006). Furthermore, risk-averse farmers may choose not to pursue such practices and potential profits, owing to the complexity and uncertainty of impacts on their particular farm operation (Bosch & Pease, 2000). In the face of skill barriers and uncertainty, group learning through a catchment community-based approach to water quality may offer a promising pathway towards sustainable water management (Dodd, Wilcock, & Parminter, 2009). Indeed, there are a number of distinct advantages to building social capital through a common property approach, particularly in promoting long-term cultural sustainability (Aminova & Abdullayev, 2009; Burton & Paragahawewa, 2011).

1.4 Motivation

There are two particularly troubling aspects for the potential success of collective approaches to managing nitrate losses, compared to managing other aspects of water quality or water quantity. First, the economics of nitrate management are characterised by strong incentives for defection, unlike other aspects of water quality addressed through more simple technology or management changes. Due to the relatively strong coupling of nitrate losses, land use intensity and profit, users can earn significant gains by deviating from agreed levels of nitrate loss.

Second, groups that collectively manage nitrate lack an inherently strong sanction, such as the opportunity to restrict a user's access to water in an irrigation scheme. Recent developments in the private sector suggest that sanctions could be developed to address this. Milk companies in New Zealand have branched out into audited self-management for environmental matters in recent years, including the use of rewards and sanctions. For example, Synlait offers price premiums for milk from suppliers that meet higher environmental standards (Synlait, 2016). Fonterra Cooperative Group has instituted a regime whereby continued non-compliance with standards (including environmental standards) results in the company refusing to pick up milk from a supplier (Fonterra, 2007). This is a particularly strong sanction due the high cost imposed on a farm unable to sell its milk. Accordingly, this could be expected to induce a high level of cooperation in a group management situation, but this has only been applied at the nation-wide level so far, rather than in respect to a specific commons management issue.

As there is growing interest in the potential to manage water quality through collective approaches, it is important to develop deeper understanding of the potential advantages and drawbacks through as many channels for research as possible. One particularly promising channel for this is the use of economic experiments. Experiments offer the advantage of systematically testing specific aspects of a particular issue, while still providing data based on the interactions of real people, rather than stylised models of economic actors. The use of a laboratory environment enables a useful level of control over the parameters of the environment, supporting greater levels of attribution for a particular given

treatment variable. Clear procedures enable the replication of experiments from one researcher to the next, which enables cumulative development of knowledge, as new experiments build on the results of previous experiments. In addition (and of particular interest for policy development), laboratory experiments enable hypotheses about particular policies to be tested in a relatively low-cost environment. Failed policies are expensive in both implementation costs and potentially adverse reactions to a policy. While we can never be completely sure how a given scenario will translate from the laboratory to the field, initial experiments can provide insight into potential advantages, pitfalls and key design parameters (Fisher, Wheeler & Zwick 1993; Levitt & List, 2007; Ostrom, 2006).

Collective management of nitrate losses offers a coordination problem typical of the design of many economic experiments concerned with public goods and common pool resources. However, there is a critical gap in our understanding of how strong, non-linear incentives for defection (as found in the nitrate management problem) interact with sanctioning institutions within a decentralised management structure. Previous research in the area has found sanctioning institutions to be relatively ineffective in this type of incentive environment in the absence of communication (Cason & Gangadharan, 2015; Cason & Gangadharan, 2016; Ostrom, Walker, & Gardner, 1992). This is potentially problematic, as in the field, communication can be expected to break down periodically. Members of the group will change over time. Expectations and norms may not be clearly understood by new members. People may enter groups that do not share the strong social capital that built the group in the beginning. In these cases, communication may not be fully effective and free-riding (the capture of benefits that have not been paid for) may be attempted. If this occurs, there is a risk that other members of a group may follow a free-rider and cooperation will break down, to the detriment of all group members and the environment. Strong institutions are needed to overcome these periodic disturbances and ensure that other members of a group continue to see value in collective management, even when some individuals do not play by the rules.

1.5 Aims

Experimental economics literature has provided many insights into the influence of norms, communication and monitoring on cooperation in social dilemma environments. However, the findings in relation to sanctions in self-managing groups are mixed across different incentive environments. In common-pool resource games, sanctions administered by group members have been found to be only marginally effective without communication (Cason & Gangadharan, 2015; Cason & Gangadharan, 2016; Ostrom et al., 1992). It is possible that the complexity of these games causes confusion, leading to participants not understanding why they are being punished. However, a systematic study of the influences of the effectiveness of sanctions and the incentive environment has not yet been undertaken.

The central aim of this study is to determine whether strong sanctions administered by members of a group are able to improve cooperation in a collective management scenario where there are strong incentives to defect, no communication and complexity makes optimisation challenging. Economic experiments are undertaken within an incentive environment designed to make a corollary for the dilemma faced by farmers engaged in collective management of nitrate loads to a water body. By further developing our understanding of the relationship between incentives and sanctions, we hope to identify areas for focus by policy-makers in developing institutions to support collective management, ensuring that they are set up for success rather than failure.

2. Literature Review

2.1 Non-point source pollution and common pool resources

Non-point source pollution is sometimes referred to as a common pool resource problem. Common pool resources are typically renewable environmental goods that provide additional benefits with increased use, up to a point where the resource becomes degraded and benefits begin to decline (Hardin, 1968). This benefit-maximising level of use is typically referred to as maximum sustainable yield. Examples of common-pool resources include forests, fisheries and water resources used for irrigation. In the context of pollution, the common resource is the capacity of a water body to assimilate contaminants. However, there is one often critically-overlooked difference in this case. One of the key conditions for successful CPR management is the salience of the resource condition to the livelihood of its users (Ostrom, 2008). For example, if an irrigation resource becomes overused, then all users are negatively affected by reduced irrigation reliability. This creates a strong motivation for the development of institutional arrangements to manage the resource effectively. In the case of pollution, the degradation of the resource (assimilative capacity) does not affect the ability of users to provide for their livelihoods (through activities like farming or cropping). This means there is no direct inherent motivation for users to manage the resource better, thus complicating the management of the issue relative to where any detrimental effects of common pool management are experienced directly by the resource users. This being said, external pressures from the broader community such as the threat of regulation have been known to drive collective action (Bay of Plenty Regional Council, 2012; New Zealand Landcare Trust, 2012). In any case, non-point source pollution shares many other characteristics of common pool resources, particularly in relation to the factors that affect management. If costs are imposed on users for the degradation of the resource, the incentive structure becomes identical.

Costs associated with the effects of non-point source pollution might be sourced endogenously or exogenously. In the first case, the source of cost for producers of non-point pollution is endogenous damage experienced directly as a result of their

actions. For example, overall levels of nitrate leached from farming operations may breach drinking water standards for aquifers that farmers depend on for their own drinking water supplies. There are few examples of this occurring in the field in a way that introduces sufficient incentives to stimulate collective action. In the second case, costs might be due to some form of group tax imposed by a regulator, as suggested by Segerson (1988). In this case, failure to protect the resource as a group could be met by sanctions against all group members by the regulator. This is problematic in reality, since producer benefits are unknown and costing environmental damage is difficult, making it difficult to determine the optimal tax level (Suter et al., 2008). Nonetheless, a few examples of this approach to regulation now exist in the field. For example, under the Everglades Forever Act in the United States, the government instituted a cropland tax based on aggregate phosphorus contamination from agricultural runoff (Athanassoglou, 2010). There is no evidence of the effectiveness of this programme however (certainly not within the academic literature).

There are also instances in which the *threat* of more severe regulation has inspired collective action, where desired community environmental outcomes are not met (Kingi, Park, & Scarsbrook, 2012). This has been suggested as a potentially productive role for the state in supporting voluntary action (Segerson & Wu, 2006). In this context, the incentive for protecting the resource effectively reflects the increasing risk of costly regulation experienced by individual firms. This introduces the salience condition, effectively turning the non-point source pollution problem into a common pool resource problem. Non-point source pollution framed as a commons problem in this way is likely to be subject to many of the limitations of common pool resource management found through field research (Cox, Arnold, & Villamayor-Tomás, 2010).

For a common pool resource, even though an individual's extraction of the resource impacts on that of others, it is difficult to exclude users. That is, the resource is *rival* but not *excludable.* This explains some of the difficulty in managing common pool resources effectively, when examining the nature of the property rights associated with this resource. For private goods (individual property), the ability of one user to exclude another's use provides the potential

for the overall level of use to be managed, provided there is some means of securing private property against theft. Club goods are excludable, but not rival. This means that artificial scarcity may be created, but there are not significant implications for the degradation of a resource in the absence of the club. For public goods, one individual's extraction has no impact on others, but due to the lack of excludability, under-provision is common if payment for the public good cannot be effectively coordinated. Common pool resources present a similar challenge to public goods, in that coordination is required to reach a socially-optimal outcome. Instead of under-provision as in the public good case, the concern is over-extraction of a common pool resource by an individual impacting negatively on other users (and indeed, potentially creating negative externalities if the resource has multiple environmental services). Common pool resource and public goods have unique problems due to the lack of excludability. The characteristics of these different types of goods are outline in table 1 below.

Table 1: Excludability and rivalry for different types of property

Property type	Excludable	Rival
Private goods	Yes	Yes
Common pool resources	No	Yes
Public goods	No	No
Club goods	Yes	No

The potential problems associated with unmanaged common pool resources were initially highlighted by Gordon (1954) and later Hardin (1968) in his seminal work, *The Tragedy of the Commons*. Gordon and Hardin independently argued that CPR users are inevitability driven to over-exploit the resource (Gordon, 1954; Hardin, 1968). However, the central thesis of their theoretical analyses is now generally accepted to be incorrect due to analysis of successful collective management of common property in the field (Feeny, Berkes, McCay, & Acheson, 1990). The problem so richly described in these works is perhaps more accurately referred to as the tragedy of open access, where common pool resources are not managed by any particular institutional framework.

Communities of resource users have frequently devised group management approaches throughout history. Indeed, Hardin himself used the example of a common pasture being over-exploited by individual graziers, which was an example of successful management in reality. Historically, common pastures were governed by communities of graziers with clear rules in relation to the number of cattle that were allowed to be grazed by any individual (Feeny, Berkes, McCay, & Acheson, 1990).

As Hardin contended that common pool resources represented a problem of property rights, his conclusion called for privatisation or nationalisation in order to control access. Research into community resource management institutions has presented a far more nuanced view (Cudney-Bueno & Basurto, 2009). There are numerous examples globally of the management of common pool resources using institutional solutions based on community organisation and collective management, rather than property rights (Burger & Gochfeld, 1998). The group property arrangement completes the potential management approaches outlined in table 2 below. In some cases, common property approaches may be more effective than nationalisation or privatisation in avoiding over-extraction (Adger & Luttrell, 2000; Ostrom, 2006), but the relationship between communities, markets and states is complex and evolving (Bowles & Gintis, 2002b). Accordingly, a rich body of research continues to emerge around effective CPR management, community governance and institutions.

Table 2: Property rights regimes in relation to resource management

Property Rights	Characteristics
Open access	Absence of enforced property rights
Group property	Resource rights held by a group of users who can exclude others
Individual property	Resource rights held by individuals (or firms) who can exclude others
Government property	Resource rights held by a Government that can regulate or subsidise use

(Ostrom, Burger, Field, Norgaard, & Policansky, 1999)

Theoretical analyses of the problem of public goods arose around the same time as the theoretical assessments of the common-pool resource problem (Ledyard & Roberts, 1974; Samuelson, 1954). Public goods differ markedly from CPR due to the lack of rivalry, as one person's benefit from a public good does not impact on those experienced by another. However, there are similarities in terms of the social nature of the problem. At their heart, both common pool resources and public goods present society with a tension between what is best for an individual, and what is best for society as a whole. Both are social dilemmas that create a risk of selfish individuals free-riding on the positive contributions of others. In experiments simulating these economic environments, some similar patterns of behaviour emerge. Accordingly, while the specific focus of this study is a common-pool resource, analysed through experimentation, there are also important lessons from the literature on public goods experiments which are considered in this review.

2.2 Experimental economics and environmental policy

Much public policy of the last century has been built on neo-classical models of rational self-interested actors. This policy approach dates back to the writings of David Hume, who considered that if policy were written for "knaves", regulators would at least do no harm (Le Grand, 1997). This position stems from a fundamental theoretical debate, as to whether humans are en-masse, fundamentally altruistic (concerned for the welfare of others) or free-riding (concerned only with their own welfare). This question is relevant to policy in the identification of the conditions necessary for commons to be protected, or for public goods to be provided. Hume's conjecture was that public policy that treated all citizens as potentially free-riding would be best, and at worst do no harm (Hume, 1875).

More recent research has shown this assumption to be incorrect. In *Covenants with and without a sword*, Ostrom, Walker and Gardner (1992) challenged the fundamental assumptions behind the Hobbsian view that sovereigns must be created in order to control subjects and administer sanctions. Their experiments

found that the highest welfare was generated through internal covenants with freely-chosen sanctioning institutions. Even more illuminatingly, covenants with no sanctions were found to be more effective than sanctions without covenants.

The treatment of diverse people as self-interested, as suggested by Hume, can undermine intrinsic motivations. In some cases, policy based on the neoclassical model of self-interested actors can even cause civic-minded individuals to behave selfishly (Bowles, 2008). Policy based on the assumption of free-riding often utilises formal fines, sanctions or incentives. In practice, there is no single rule for how fines or incentives will affect intrinsic motivations, such as altruism or ethical norms. Policies based on formal sanctions or incentives may conflict with or complement intrinsic motivations, depending on how they are implemented (Bowles & Polanía-Reyes, 2012). Policies designed with behavioural nuances and heterogeneity in mind can avoid adverse reactions to incentives and induce higher levels of efficiency. It is thus critically important to understand when incentives (in the neoclassical sense) do and do not work (Gneezy, Meier, & Rey-Biel, 2011).

One of the most common methods employed by a society to induce more cooperative behaviour is the use of sanctions against those who fail to comply with a social norm (an informal understanding governing social behaviour). Formal sanctions tend to be more tangible (such as fines or incarceration). Informal sanctions may be as simple as the expression of disapproval. Formal and informal sanctions may operate as complements or substitutes, depending on the arrangement. Blau (1964) theorised that legitimate power structures develop as an institutionalised means of upholding social norms. However, formalised institutions also react with and affect social norms. In some cases, a law may crowd out an existing social norm, substituting state control for a less formal institution and reducing efficiency (Bohnet, Frey, & Huck, 2001; Coleman, 1993; Fehr & Gächter, 2001; Frey, 1997). Alternatively, formal sanctions may correct inefficient social norms (Posner, 1996). In other cases, laws or contracts may provide signalling that supports a social norm and vice-versa, resulting in higher net social benefits due to the low administration burden. In these environments, formal and informal sanctions in relation to norms act as complements (Kube &

Traxler, 2011; Lazzarini, Miller, & Zenger, 2004; Noussair & Tucker, 2005; Posner, 1997). These interactions of norms, laws and sanctions (both formal and informal) are of critical interest to policy-makers, as there are profound implications for the costs of implementation, monitoring and enforcement.

Sustaining cooperation is critical to many facets of governance in human society, including the management of natural resources. Common pool resources in particular require high levels of cooperation in order to define appropriation and provision rules, to ensure common property is not over-exploited. Many of the most promising insights have been provided through experimental studies. Institutional structures and incentives are now being systematically tested for their potential for sustaining cooperation in providing public goods and managing common pool resources. Economic experiments have the advantage of providing significant control over the conditions likely to affect behaviour, supporting designs which can begin to tease apart the complex interactions of factors that affect decision-making. Good experiments enable different incentive environments to be tested, to see whether participants behave in a manner consistent with economic theory or whether there are other forces in play. The findings from such experiments provide useful behavioural insights into how policies may need to be adapted in order to be effective or efficient in the field (Ostrom, 2006; Vossler, Poe, Schulze, & Segerson, 2006).

The choices made by participants in social dilemma games are typically described in relation to two different strategic equilibria, the Nash equilibrium and the social (Pareto) optimum. These two arrangements represent a non-cooperative and a cooperative equilibrium, respectively. The Nash equilibrium occurs where all participants seek to maximise their own returns, in the absence of any knowledge about others' strategies (Nash, 1951). This is comparable to what might be expected under neo-classical economic theory. Each individual assumes that all other individuals are similarly seeking to maximise their individual welfare, rather than the interests of the group as a whole.

The social optimum occurs where all participants seek to maximise the total returns to the group, and assume that all other participants are trying to do the

same. This produces a solution to the game which is Pareto-optimal and maximises total welfare. In the neo-classical model, it would be expected that some form of intervention would be required in order to achieve such a social optimum. Early in the experimental fields of sociology, psychology, political science and economics, free-riding was found to be less common than expected and explanations in terms of "rational" actors elusive. Instead, a diverse range of behaviours and strategies were uncovered, ranging in between these two extreme equilibria (Bohm, 1972; Isaac, Walker, & Thomas, 1984; Marwell & Ames, 1979; van de Kragt, Dawes, & Orbell, 1988).

In one-shot experiments that do not feature institutional controls, cooperation usually occurs at a level higher than the self-interested Nash equilibrium, but less than the social optimum, with contributions to a public good typically at 40 to 60 per cent of the initial endowment. This was a key finding of early public goods experiments (Dawes, McTavish, & Shaklee, 1977; Marwell & Ames, 1979, 1981) but was found to be unstable in repeated one-shot games. With repetition of a game, while contributions generally begin at a similar level between the social optimum and Nash equilibrium, they tend to decline over time towards the Nash equilibrium (Andreoni, 1988; Isaac, McCue, & Plott, 1985; Isaac & Walker, 1988b).

A number of factors influencing the level of cooperation over time have been demonstrated through systematic experimentation. One of the early factors found to improve the rate of contributions is communication (Dawes et al., 1977; Isaac & Walker, 1988a; Ostrom et al., 1992), which is discussed further below. More recently, two major advances have been made – the existence of "conditional co-operators" and the role of costly punishment (Chaudhuri, 2011).

While some participants follow self-interested strategies, others behave randomly or altruistically. The largest group, however, tends to be conditional co-operators, who choose their strategies based on expectations and evidence of others' behaviour (Fischbacher, Gächter, & Fehr, 2001). Whether exploitative or fairness-minded behaviour dominates is dependent on the economic environment (Fehr & Schmidt, 1999). Conditional co-operators tend to contribute according to

others' behaviour. If they find that they have contributed less than the group average, they contribute more in the next round. If they find that they have contributed more than the group average, they will tend to contribute less (Bowles & Gintis, 2002a; Fischbacher et al., 2001). The commonly-observed decline in contributions over time can be explained by the fact that conditional co-operators tend to have a bias towards contributing slightly less than the average. This means that most participants will follow the declining average downward over time (Fischbacher et al., 2001). This makes sustaining cooperation very unlikely in unregulated institutions.

Experimental work has led to the description of participants in experiments in a number of different ways, according to their general behaviour in response to a given economic environment. This provides a more complete description of likely behavioural responses to a given policy scenario or decision-making environment. For example, a five-agent model used by Carpenter involved unconditional co-operators (altruists), unconditional free-riders (knaves), general tit-for-taters (conditional co-operators), nice punishers (unconditional co-operators that also punish free-riders) and mean punishers (essentially hypocrites that both free-ride and punish). The model essentially captures the three different cooperation preferences, as well as attitudes towards the use of sanctions (Carpenter, 2007b).

2.3 Social dilemma game environments

There are many experimental games used to measure social preferences (for example, the Prisoner's Dilemma game, the Ultimatum game or the Dictator game). The class of games most pertinent to this research (and thus the predominant focus of this literature review) is broadly referred to as public goods games, though this also includes common pool resource games (Ledyard, 1997). These games involve multiple (usually four or more) players faced with a social dilemma, where the interests of individuals diverge from the social optimum. Participants must cooperate in order to gain the maximum social benefit.

While some experimental findings are relatively consistent, other findings vary markedly across the different types of games, such as those associated with the overall level of cooperation, the degree of efficiency and the effectiveness of sanctions (Cason & Gangadharan, 2015). This variation highlights the importance of considering the experimental context, incentives and institutional structure when interpreting experimental results. While sharing a divergence between individual incentives and the promotion of welfare maximisation, there are critical differences in the construction of payoff functions, rivalry and institutions that have consequential impacts on the incentive structures within each game. This means that while some findings are common, we cannot generalise between the different types of game. There are four key types of public goods games of relevance to this research. A brief description of key experimental settings follows.

Linear Voluntary Contribution Mechanism (VCM)

The linear VCM is perhaps the most frequently-used of all public goods games. The linear VCM operates by assigning all members of a group an initial endowment of tokens (E). Participants can then choose to contribute some or all of their tokens (g) to a public account, in order to generate a given rate of return (β). Tokens retained in the participant's private account receive the rate of return (α). All participants then receive equal shares of the public account revenue. An individual's payoff (π_i) can thus be calculated by equation 1.

$$\boldsymbol{\pi_i = \alpha(E_i - g_i) + \beta \sum_{j=1}^{n} g_j} \qquad \textbf{(1)}$$

This equation is typically parameterised so that $\beta < \alpha$, which makes the neoclassical self-interested (Nash) strategy for an individual to be retention of all their tokens. Participants that retain all their tokens can still get a return from the group account by free-riding on the contributions of others. In addition, the game is typically parameterised so that $\alpha < n\beta$ (where n denotes number of participants), which makes the social optimum involve all participants contributing all of their tokens. The Nash and social optimum strategies are thus at opposite, extreme

ends of the decision space (g). This presents a simple stylised version of the public goods problem. The whole group benefits from the provision of the public good. However, as no individual can be excluded from those benefits, there is no self-interested incentive to pay for it. This is true of typical public goods, such as defence.

Non-Linear VCM

The non-linear VCM is very similar to the linear VCM, except that additional complexity is introduced by varying the payoff parameters α and β so that the marginal per-capita group return decreases relative to the individual return as contributions increase. The result is a much more difficult game, where the Nash equilibrium and social optimum may be in the interior of the decision space (g) rather than at the level of maximum or minimum contributions. This is similar in terms of the social dilemma that is described by the linear VCM, in that individuals are not incentivised to contribute a sufficient amount to the public good to reach the social optimum level. The variation that makes for the interior Nash equilibrium and social optimum makes it more realistic. Continuing with the example of defence, it is unlikely that an individual would wish to pay nothing and be entirely unprotected, but it is equally unlikely that contributing the maximum possible amount to defence would be worthwhile. Above a certain level of defence, additional spending would return only minimal gains, which would not be worthwhile for society. In the field, many public goods are subject to diminishing marginal returns in this way.

Common Pool Resource (CPR) Game

CPR games are an even more complex non-linear game, designed to emulate a renewable resource management environment, with a maximum sustainable yield. The explanation of the experimental arrangement differs slightly in that each token "invested" in the public account corresponds to a level of extraction from a common pool resource (for example, a forest or fishery). Participants may choose to place tokens in a private or public account. For the example of a forest, the public account represents effort spent harvesting wood, while the private account represents effort on the next best alternative. The return to the group increases

with greater harvesting of the forest up to a point of maximum sustainable yield. If too much harvesting occurs (i.e. too many tokens are placed in the group account), the group returns become worse, as not enough new trees are maturing to replace the ones that have been harvested. Importantly, returns to each individual depend on how much effort they expend, rather than an equal share. While the social optimum is to only harvest the maximum sustainable yield, each individual can improve their own returns by harvesting more than their share of wood, even if this means degrading the resource. This sets up the twin problems of non-excludability and rivalry that define common pool resource problems in the field.

The following description takes the form used by Ostrom, Walker and Gardner in their seminal paper *Covenants with and without a Sword* (1992). For ease of comparison, the same labels are used as for the VCM description above. Participants can choose to invest some or all of their tokens (g) from their initial endowment (E). Group payoffs are determined by a concave function (F). This function is parameterised so that $F(0)=0$, $F'(0)>\alpha$, and $F'(nE)<0$. This means that both the Pareto optimum and Nash equilibrium are internal to the extraction choice space (g). The Pareto optimum occurs at slightly less than the maximum productivity of the CPR. The Nash equilibrium results in over-exploitation, but not exhaustion of the CPR.

$$\boldsymbol{\pi_i = \alpha(E_i - g_i) + \frac{g_i}{\sum_{j=1}^{n} g_j} F \sum_{j=1}^{n} g_j} \qquad \textbf{(2)}$$

Importantly, in this structure individual payoffs from the group account are earned according to the level of appropriation (determined by $g_i / \sum g_i$) rather than a pre-determined share ($1/n$) as in VCM games. This introduces rivalry to the game, so that when one participant increases their own contribution, their appropriation gains them much higher benefits, while decreasing the earnings of others.

Pollution games and ambient schemes

Pollution games and ambient schemes are similar to CPR games, in that they tend to have interior Nash and Pareto solutions. Payoffs in relation to use of the CPR are also associated with a negative impact driven by the sum of group contributions. The differences are primarily driven by parameterisation and framing. These pollution games were structured to test experimentally the theoretical work of Segerson and others on "ambient schemes" for pollution control.

A typical formulation of this game is used by Cason and Gangadharan (2013), with payoffs and costs structured similarly to Spraggon (2002). Individual payoffs are expressed according to a concave function, so that the benefits (*B*) for a given level of emissions (*E*) are expressed as shown in equation 3.

$$\boldsymbol{B(E_i) = 44 - 0.002(125 - E_i)^2} \qquad \textbf{(3)}$$

In this example, maximum benefits of 44 are achieved at an emissions level of 125. The concavity of the function reflects the fact that above a certain level, excessive use of inputs causes decreased benefits, reflecting declining marginal returns for the individual. Participants choose a given level of emissions and receive the associated level of benefit.

Environmental damages (*ED*) are expressed by way of a damage function. Each participant experiences an equal (1/*n*) share of the cost of these damages, calibrated by a transmission coefficient (*t*). Thus, the return to each individual is expressed as $\pi_i = B(E_i) - ED$, where $B(E_i)$ is defined in equation (3) and *ED* is defined in equation 4 below.

$$\boldsymbol{ED = \frac{1}{n}\left[\sum_{j=1}^{n_j} tE_j\right]} \qquad \textbf{(4)}$$

There are two main ways in which users might experience damage costs in the field. The first is if they are imposed exogenously on all members of the

population. For example, a regulator might impose a tax on the diffuse emissions of nutrients that are measured in a water body, as suggested in ambient scheme institutions (such as Spraggon, (2002) for example). The second formulation is to assume that damages are experienced endogenously (i.e. by the producers themselves). The example used by Cason and Gangadharan is that of pesticide leaching, which might pollute the drinking water supplies of individual farmers (Cason & Gangadharan, 2013).

Imposing this damage function makes the incentive environment for pollution games very similar to that of a common-pool resource. Since each unit of production for an individual adds to the damage function experienced by the group, subjects are able to externalise most of the damage they cause while earning significant benefits for themselves, making the game rival. Though the public and private benefit (or damage) functions are different to those in CPR games, individuals experience a similar concave total benefit function, with their benefits being rival with other participants in the game.

2.4 Incentives – group size, marginal per-capita return and rivalry

The design of incentive environments has a strong impact on the overall level of cooperation within a game. Accordingly, incentives must be included in consideration of how other institutional factors, such as sanctions, affect behaviour. This is true within the same type of game (for example, where payoff functions or marginal per capital return may vary) and across different games (such as the rival and non-rival designs of CPR and VCM games, respectively).

Early in the development of public goods theory generally, the size of groups involved in social dilemmas was theorised to be of critical influence on whether or not cooperation is likely to occur. Olson's (1965) seminal work on collective action paid particular attention to how in large groups, the incentives for the provision of public goods decline, as each individual would gain a smaller and smaller share of the benefits as group size increased. However, this conception of the group size effect is inherently intertwined with the issue of incentives.

Marwell and Ames (1979) showed early on that there was little effect from group size itself, but a strong effect from the marginal per capita return (MPCR). Unless experiments are designed in such a way as to alter the MPCR in concert with group size, the observed group size effect is due to the changing MPCR. The exception to this is that in the field and in experiments where mutual monitoring is required, group size is likely to have effects on the effectiveness of monitoring and enforcement (Carpenter, 2007b). The effectiveness of monitoring is also covered extensively in the field literature on CPRs. For example, both the work of Ostrom (1994) on design principles for common pool resource governance and subsequent reviews such as Cox, Arnold and Villamayor-Tomás (2010) emphasise the importance of effective monitoring for effective governance.

Isaac and Walker (1988b) and Isaac et al. (1984) both investigated changes to the MPCR in terms of its effect on cooperation. This work identified that a higher MPCR was found to increase the rate of cooperation, independent of group size. Increasing the MPCR changes the incentive environment by increasing the benefits of cooperation, relative to free-riding. Interestingly, while less free-riding was found to occur in these experiments than would be predicted by the classical model, participants still responded clearly to incentives. While a significant amount of work on marginal incentives has been undertaken in linear VCM games, there has been little systematic examination of the effect of incentives across different game environments. Cason and Gangadharan (2015) undertook some recent work in peer punishment comparing the linear VCM, non-linear VCM and CPR environments while holding total payoffs constant, but this work did not take into account the significantly different marginal incentives to defect created by the rival and nature of the CPR game.

A third critical impact on the incentive environment in social dilemma games is whether there is the presence of rivalry. Rivalry has long been theorised as a critical difference between public goods and common pool resources (Musgrave, 1959). However, it has not always been recognised as such in the experimental literature (for example, Ledyard, (1997)). In fact the difference is marked. In the context of a public good such as defence, one person's benefit from the good does not impact on another's. However, for a common-pool resource like a forest,

once wood has been harvested, it is no longer available for others to use. VCM games assign benefits from the group account to each participant based on a pre-determined formula, so that all participants' gains are dependent only on the total level of the good provided, rather than their own individual contributions. For CPR games, however, benefits from the public account are assigned according to the level of participants' contributions. Systematic investigation of this difference was undertaken by Apesteguia and Maier-Rigaud (2006) in experiments featuring identical payoffs and social optima. This study identified significant differences in the distributions of how participants respond to the two different rival and non-rival environments, though the general pattern of beginning close to the social optimum and declining towards Nash equilibrium was still present.

2.5 Sanctions

The use of sanctions in sustaining cooperation is central to the aim of this study. Fundamentally, we are attempting to determine whether strong incentives for free-riding can be overcome by strong sanctions. The use of sanctioning institutions to improve provision of public goods in experiments is extensive and has been known for some time (Yamagishi (1986) for example). Sanctions may be administered by group members to each other (peer punishment), or by an external sanctioning institution (simulating a government intervention). Punishment may be formal (such as monetary fines), or informal (such as social disapproval). Formal and informal sanctions can work independently or in concert to enhance welfare (Dugar, 2013; Noussair & Tucker, 2005). There are two ways in which sanctions can enhance cooperation. Directly, sanctions can change the payoffs of cooperation or defection in such a way as to make cooperation preferable even to free-riders. Indirectly, sanctions promote cooperation among conditional co-operators by ensuring that they are not discouraged by free-riders (Shinada & Yamagishi, 2007).

Peer punishment is perhaps the most studied of sanctioning mechanisms, particularly in recent years. Typically, when peer punishment is employed within a given experimental setting, participants are able to pay a small fee and

consequently lower the earnings of others by a larger amount. This may be unconstrained (where any participant can punish any other, regardless of their choice of cooperation or defection) or subject to constraints (for example, through restricting punishment to fining defectors only) (Casari, 2005). This type of arrangement reflects more informal community-based governance approaches in the field, recognising that sanctioning another member of a group is never without cost, whether this cost is borne in terms of time, effort or social capital (Casari & Plott, 2003). Alternatively, experiments with some kind of centralised sanctioning institution typically involve standardising the punishment of individuals according to a set of rules, sometimes with rules designed by participants themselves. This bears closer resemblance to the delegation of enforcement to state power and the creation of a Leviathan (i.e. a large central body that is in charge of regulation). Peer punishment is common in VCM and CPR game settings, while state punishment has dominated pollution game settings – though both approaches have been used in all game environments to some extent. The reason for this apparent partitioning may be that the VCM and CPR game literature is primarily focused on group dynamics and behaviour in mitigating the social dilemma, while pollution game literature has been more focused on the overall effectiveness of the institution in achieving an environmental outcome in circumstances where the state cannot observe the behaviour of individual firms.

Many explanations have been advanced for people's demand for the use of peer punishment. Some experiments suggest that (at least in laboratory settings) while strategic motivations drive the allocation of rewards, more intrinsic motivations drive punishment (Choi & Ahn, 2013). Notably, the use of punishment is common even at the end of a game when there is no subsequent benefit, indicating that participants enjoy punishing other participants for reasons other than merely influencing the actions of others in a game (Fudenberg & Pathak, 2010). Other possible explanations include inequality-aversion, emotions, reciprocity, confusion, spite and social norms. These motivations are difficult to discern and while theories have been advanced, this remains an open empirical question (Casari, 2005; Falk, Fehr, & Fischbacher, 2005).

Peer punishment has been particularly studied in recent years in linear VCM games, especially in relation to its capacity to improve cooperation[1]. Fehr and Gächter (2000) generated significant interest with their findings in relation to costly punishment, triggering a wave of research under similar designs and punishment technologies. Their key findings were:

1. That there is significant willingness of co-operators to punish free-riders.
2. That this holds true even when punishment is costly to the punisher.
3. That free riders are punished more heavily, the more they deviate from the contributions of co-operators.

This phenomenon is sometimes referred to as "altruistic punishment", due to the fact that punishers are not expected to gain any benefits from the money they spend on costly punishment (Fehr & Gächter, 2002). Because of this cost, punishment is itself a public good. The use of punishment increases pro-social behaviour, and is therefore welfare enhancing. However, this welfare enhancement is non-rival (since all group members benefit from the punishment of free riders) and non-excludable (since there is no means to limit these benefits to the punisher alone). The use of costly punishment by Fehr and Gächter (2002) increased players' contributions significantly and prevented the typical decline in contributions observed over time in experiments without punishment. Significantly, this pattern occurred even with designs based around strangers and with no communication.

Falk et al. (2005) undertook subsequent work examining the motivations behind the use of costly punishment. Interestingly, while most punishment undertaken by co-operators was targeted towards defectors, defectors also undertook a significant amount of punishment towards co-operators. By manipulating the degree to which participants could affect the difference between each other's payoffs, the authors were also able to determine that punishment was largely driven by the desire to punish wrongdoing than by a desire to equalise benefits (Falk et al., 2005).

[1] For a survey, see Chaudhuri (2011).

Experimental studies giving participants the option of joining sanctioning or non-sanctioning institutions have demonstrated that most participants prefer to have some form of recourse to punish free-riders. Gürerk, Irlenbusch, and Rockenbach (2006) provided an illustrative example of the advantage of sanctions. In an experiment giving participants the choice of whether to join a sanctioning or non-sanctioning institution, over time the non-sanctioning institution became completely depopulated due to the competitive advantage created by sanctions (Gürerk, Irlenbusch, & Rockenbach, 2006). Whether sanctions improve efficiency (and enhance overall welfare compared to not having punishment, after the costs of punishment are accounted for) is not straightforward. The impact of punishment on efficiency depends on the incentive environment and the cost of punishment, as well as the time horizon (Chaudhuri 2011).

Gächter, Renner and Sefton (2008) have comprehensively shown sanctioning institutions to be more efficient in linear VCM games, given a large enough number of periods. Over time, participants learn to centre contributions around norms which do not attract punishment, which makes punishment less necessary. While contributions remain high, less punishment is needed in later rounds of the game. The initial costs of punishment in establishing the norm become insignificant as earnings continue to be high in subsequent periods with little or no punishment (Gächter, Renner, & Sefton, 2008). Combining informal sanctions with peer punishment has also been shown to enhance welfare. As informal sanctions support greater cooperation without associated costs, the combination of both formal and informal sanctions results in higher efficiency than games featuring only formal fines (Dugar, 2013; Noussair & Tucker, 2005).

Though much of this recent work has been undertaken in VCM public goods experiments, there is a parallel history of investigation into punishment in CPR experiments, though the examples are much less frequent. This is surprising, given that peer punishment is frequently found in the field as part of institutions that have been successful in sustaining common-pool resources (Cox, Arnold, & Villamayor-Tomás, 2010; Wilson, Ostrom, & Cox, 2013). Many studies of diverse resource user groups around the world have found sanctions to be a key part of sustaining cooperation over time in successful institutions (Gibson,

Williams, & Ostrom, 2005). Ostrom, Walker and Gardner's (1992) seminal work in resource management institutions examined the role of peer punishment, alongside less formal covenant arrangements, using CPR games. Interestingly, this study reached the opposite conclusion to that of Fehr and Gächter (2000). Namely, that in the absence of communication, peer punishment did not improve welfare and in some cases, was even worse than unregulated games (Ostrom, Walker, & Gardner, 1992). The potential causes of these divergent findings have received little attention until recent years. Kingsley (2015) tackled the issue of the different incentives to defect by testing peer punishment in two environments (a public goods game and a CPR game) with identical payoff curves, Nash equilibria and social optima. The effectiveness of peer punishment in the two environments was, in fact, very similar. This suggests it is entirely possible that responsiveness to incentives or punishment is identical across the different environments, but that that the welfare-enhancing results of Fehr and Gächter (2000) are possible only in the simple VCM environment.

Sanctions have also been studied extensively in pollution games, usually through the use of centralised formal sanctioning institutions, though a few examples of peer punishment have been investigated (Cason & Gangadharan, 2013). Though the experiments have a very similar incentive environment to common-pool resource games, the literature is worth examining in some detail due to the careful consideration and definition of the non-point source pollution problem that has been developed and the accompanying application of sanctioning mechanisms.

Considerable attention has been paid in the experimental literature on non-point source pollution to so-called "ambient schemes". Ambient schemes were devised as a theoretical response to the problems of uncertainty and unobservability associated with NPS pollution. Segerson (1988) designed one of the first such systems as a means to apply environmental taxes where emissions are stochastic and not directly observable, building on the work of Holmstrom (1982) in the moral hazard literature, as related to labour economics (Holmstrom, 1982; Hölmstrom, 1979; K. Segerson, 1988). Segerson's ambient schemes attempted to resolve the stochasticity problem by applying taxes and rebates based on the amount of a pollutant observed in the environment (for example, sediment in a

river) following the logic that polluters would adjust their mitigation behaviour so that the probability distribution for emissions shifted towards the desired equilibrium outcome. The observability problem was intended to be resolved by applying taxes to all possible polluters at the same rate, equal to the marginal abatement cost (i.e. if marginal damages are estimated at $100, all polluters face a tax of $100 each for a breach of an environmental limit). Similar instruments were developed on a theoretical basis in order to address various shortcomings in relation to the incentive structures and information demands of the institutions described by Segerson (Cabe & Herriges, 1992; Hansen, 1998; Horan, Shortle, & Abler, 1998, 2002; Karp, 2005). Xepapadeas (1995) extended these instruments to involve polluting firms that could volunteer more accurate information in exchange for a lower tax rate, thus transforming the ambient taxes into dynamic Pigouvian taxes.

The first experimental work testing these institutions was undertaken by Spraggon (2002). Spraggon (2002) found Segerson's institutions to successfully mitigate the moral hazard problem regardless of uncertainty and experience, through the use of centralised punishment. However, he recognised that as none of the instruments ensure individual compliance, equity concerns were likely to be raised as members of the group that cooperate are penalised for the behaviour of free-riders (Spraggon, 2002). The work of Spraggon (2002) has since been extended to include cases where polluters might cooperate in order to game the subsidy system (Millock & Salanie, 2005; Poe, Schulze, Segerson, Suter, & Vossler, 2004; Suter et al., 2008) and to compare the relative efficiency of mechanisms including marginal taxes and subsidies, fixed penalty mechanisms, and combined fixed penalty and tax/subsidy mechanisms (Cochard, Willinger, & Xepapadeas, 2005; Francisco, Requate, & Schram, 2004; Suter et al., 2008). In each case, incremental changes have been suggested to improve theoretical shortcomings.

However, there are more fundamental practical problems associated with ambient schemes which are not addressed. Though a few examples exist, little is known about their actual effectiveness. There remains a need to evaluate such policies as it seems the theoretical and experimental settings simply do not reflect the reality

of the problem. Ambient schemes would likely work best under four conditions: (1) a small number of polluters; (2) relatively homogeneous firms; (3) readily monitored water quality; and (4) short time lags between emissions and their contribution to ambient pollution levels (Poe et al., 2004). In reality, these conditions are unlikely to be satisfied. In relation to needing a small number of polluters, many catchments are likely to be far too large to satisfy this condition. While a regulator may theoretically divide a large catchment into more manageable units, this is likely to be constrained by the ability to attribute water quality for each management unit to the appropriate group of firms. In relation to needing relatively homogeneous firms, the heterogeneity of farms and farmers is well-documented, with important implications for the cost-effectiveness of a given environmental policy (Doole & Pannell, 2012). While monitoring ambient water quality at sufficient resolution is technically possible, at present water quality modelling is constrained by cost and capacity (NIWA 2014). The requirement for short time lags is perhaps most problematic of all. In the case of sediment inputs to water bodies, there may be a relatively short lag between runoff and appearance in a water body (especially during high-rainfall events). However, the agricultural practices that led to the emission event are not necessarily undertaken recently, as it is more likely that farm management practices are designed to spread risk in anticipation of stochastic climate variation (Romstad, 2003). In the case of dissolved phosphorus runoff, time lags are likely to be too long to enable an ambient scheme to operate effectively. For nitrogen emissions, the transmission via groundwater to a surface water body may take place on the scale of decades or more (Meals, Dressing, & Davenport, 2010). This gives current ambient concentrations of pollutants, at best, a weak relationship to current land use practices. This suggests that a centralised punishment institution based on a group tax, as used in ambient scheme experiments, is not particularly realistic compared to peer punishment, where users may more readily monitor each other.

More recently, Cason and Gangadharan (2013) tested both an ambient tax and peer punishment in isolation and in combination. Results showed that although formal group tax mechanisms (as in ambient schemes) were more effective than informal peer punishment, the most effective approach for achieving a desired

environmental outcome was to use a combination of both. This research echoes earlier findings on the combination of formal and informal sanctions in VCM games. More recently, group tax and peer punishment mechanisms have been tested using heterogeneous firms. These experiments are beginning to tease out important factors which would significantly hinder the implementation of any such schemes in real policy. For example, as soon as heterogeneity is introduced to the game, the allocation of rights and questions of equity begin to hamper overall levels of efficiency and cooperation. Attempting to implement policies for heterogeneous actors without first resolving these questions is likely to prove difficult (Cason & Gangadharan, 2013; Reuben & Riedl, 2013; Suter, Vossler, & Poe, 2009).

2.7 The cost of sanctions

Within public goods games, the cost of peer punishment is typically expressed as a fee-to-fine or punishment effectiveness ratio (i.e. how many units of punishment are meted out for one unit of cost to the punisher). The imposition of this cost in experiments is generally used to reflect the fact that in the field, the punishment of others is likely to carry costs (whether in money, time, effort, social capital or some other monetary or social cost). The term *punishment effectiveness ratio* is used in this study, though in practice this is the same as cost, since the ratio reflects the sanction associated with a given cost. There are several effects that variation in this ratio has on the use of punishment, which have knock-on effects for levels of cooperation in a group. The effectiveness of punishment interacts with the incentives to defect within a given institutional structure (Murphy, Rapoport, & Parco, 2006). This is an important component to understanding how punishment is used across different economic environments, such as in the determination of differences between the VCM and CPR games.

Several studies have found that the cost of punishment affects demand. Generally, when incentives are held constant, increasing the punishment effectiveness ratio increases the amount of punishment demanded by participants. In this way punishment is a normal good, though it is generally inelastic with

respect to income (Anderson & Putterman, 2006; Carpenter, 2007a). Due to this change in demand, the cost of punishment also has effects on the overall level of cooperation within a game. Nikiforakis and Normann (2008) undertook a particularly useful systematic study using the cost of punishment as their main treatment variable in a series of linear VCM games. With a punishment effectiveness ratio of 1:1 or 1:2, punishment in their VCM environment was found to increase contributions, but be ineffective in arresting the decline of cooperation over time and ineffective in increasing welfare relative to a treatment with no punishment. A treatment with a punishment ratio of 1:3 was required in order to increase welfare and prevent the decline in contributions over time.

Counter-intuitively, participants still show a demand for punishment even when it is not credible. Experiments suggest that overall demand for punishment is a social preference, though the total amount may vary in response to price (Casari & Luini, 2012). Fehr and Fischbacher (2004) found that uninterested third party participants were willing to punish free-riders, even when this came at a cost. Similarly, while defection may be observed with and without sanctions, the degree of defection is responsive to the size of a potential sanction (Houser, Xiao, McCabe, & Smith, 2008). Given the responsiveness to price, more severe punishment (with a higher punishment effectiveness ratio) has the potential to increase overall payoffs and efficiency, due to lower cost of enforcing cooperation and greater overall levels of cooperation (Ambrus & Greiner, 2012).

Where monitoring is fully effective, increasing the severity of punishment increases the overall group payoffs. However, in a more realistic environment that features imperfect monitoring, increasing severity of punishment has been found to initially decrease overall payoffs, even in the linear VCM, as costs rise without sufficient deterrence to increase cooperation (Ambrus & Greiner, 2012). Punishment behaviour in a group dynamic also responds to variations in the cost of punishment between individuals. In cases where participants have heterogeneous costs of punishment, tacit agreements may emerge that only the strongest will punish free-riders, increasing overall efficiency by effectively centralising punishment (Przepiorka & Diekmann, 2013). Individuals will shoulder a higher burden of the punishment of free-riders in experiments where

they are given more cost-effective punishment. Nikiforakis, Normann, and Wallace (2010) found that while players with higher-punishment effectiveness contributed similar amounts, they tended to punish more than their fellow-players with less-effective punishment. This is likely to be driven by the same relationship between cost and demand observed between experiments with homogeneous costs of punishment that vary the cost across different treatments (Nikiforakis, Normann, & Wallace, 2010).

It is possible that the tendency to pool or centralise punishment institutions within society is a response developed to lower the cost of punishment. When given the option, the centralisation or specialisation of a punishment institution or enforcer can be an attractive means for reducing the cost of punishment. In a design featuring the choice of nominating a "hired gun" to act as the enforcer (often referred to as pool-punishment), Andreoni and Gee (2012) found that participants would choose this option 70 per cent of the time. This centralised arrangement was also found to be more economically efficient than peer punishment (Andreoni & Gee, 2012). Similar results were found by Traulsen, Röhl and Milinski (2012). However, some subject groups in these experiments were successful using informal peer sanctions. As further support for the hypothesis that centralisation is driven by a desire to reduce cost, the choice of punishment option was responsive to the cost of formal third-party punishment (Markussen, Putterman, & Tyran, 2014). Zhang, Li, De Silva, Bednarik, & Sigmund (2014) found that in fact pool-punishment only emerges when there is punishment of individuals that fail to punish free-riders.

The vast majority of experiments examining the cost of punishment have been undertaken in VCM games. While the experiments of Ostrom et al. (1992) included punishment effectiveness ratios of 1:4 and 1:2, there has been little systematic exploration since. CPR games typically feature higher incentives for free-riding, compared to VCM games. The less frequent punishment observed in these experiments was found by Cason and Gangadharan (2015) to be less effective in sustaining cooperation (when using the same punishment ratio of 1:3 also used in the linear VCM games). However, experiments varying the returns in VCM games have found that in fact players punish *more* when the stakes are

higher (Carpenter, 2007b) which suggests that confusion or disagreement about norms may have a bigger role to play in the CPR games. In order to properly disentangle this issue, it is necessary to answer the question as to whether a stronger punishment ratio can better support coordination in the CPR game. That is, whether more severe punishment can be an effective mechanism in promoting cooperation where punishment is less frequent.

Cason and Gangadharan (2015) have undertaken some experiments recently that suggest support for the findings of Ostrom, Walker and Gardner (1992) in the CPR environment. Cason and Gangadharan (2015) tested the impacts of peer punishment across three different types of game (linear VCM, non-linear VCM and CPR), showing differences between institutions. One of the key changes to the design of Ostrom, Walker and Gardner (1992) was to use an updated punishment institution, including the ability to vary the level of punishment used. A punishment effectiveness ratio of 1:3 was used (compared to a range of 1:2 to 1:4 used by Ostrom et al. (1992)) and each participant could assign up to five punishment points to each other participant (compared to only one per participant in Ostrom et al. (1992)). Findings from these experiments indicated that even with the updated mechanism, peer punishment was less effective in sustaining cooperation in the more complex non-linear VCM and CPR environments than in the linear VCM. The authors have suggested that this is due to the effect of confusion. However, the interaction of different incentives for defection or cooperation with the punishment effectiveness ratio is not explicitly considered (Cason & Gangadharan, 2015).

2.8 Norms and complexity

Norms are informal understandings or agreements about what the right thing to do is in a given situation. Where there is potential for private self-interest to generate a negative impact on other members of a society, demand for some kind of governing social norm exists (Coleman, 1990). The existence of cooperation among a group faced with a social dilemma is predicated on there being some kind of shared understanding of the right thing to do. This is equally relevant to

the use of sanctions in the management of common-pool resources. In social dilemma games (and indeed in many other contexts), people make decisions about their actions based on their expectations of the behaviour of others. For some, this may mean a high contribution to the public good, potentially reflecting an expectation that others will do the same. For others, this may mean a small contribution, if the expectation is that others will free-ride (Dawes et al., 1977). The decisions made do not automatically reflect the preferences of the individuals in question; rather, the decisions reflect the expectations of individuals about the preferences of others. For example, a person with strong social tendencies may choose to not make a large contribution to a public good, if their expectation is that others will simply free ride.

In the absence of institutional controls, it is therefore to be expected that the base level of cooperation that occurs is partially a product of cultural expectations about cooperation and free-riding. Cross-cultural experimental research has shown that institutional mechanisms designed to promote cooperation such as costly punishment have mixed success, depending on existing norms. In particular, whether cultural norms of cooperation are already in existence has a strong influence on the efficiency and effectiveness of institutional arrangements. Some pools of participants may predominantly punish free-riding, while others may be relatively indiscriminate, undertaking anti-social punishment. Weak norms of civic cooperation and the weakness of the rule of law in a country are significant predictors of antisocial punishment (Herrmann, Thöni, & Gächter, 2008). Interestingly, punishment may be more severe in cohesive groups, perhaps reflecting higher expectations of norm adherence. Shinada, Yamagishi and Ohumura found that participants who are cooperative in a gift-giving game punish noncooperative in-group members more severely than they punish noncooperative out-group members (Shinada, Yamagishi, & Ohmura, 2004).

In the absence of a clear allocation rule under which punishment should be applied, considerable confusion can arise due to participants having different views on the social norm. In some cases, one preference may dominate, leading to a sustainable state where punishment is based on a single social norm. However, there is no guarantee that this will happen quickly, or at all, leading to

protracted "feuds" between different group members a particular rule (Nikiforakis, Noussair, & Wilkening, 2012).

Clarity about norms and expectations is critical to sustaining cooperation. Even when games do not feature sanctions, the ability to set expectations about the best action for the group as a whole through communication significantly increases cooperation. The use of punishment without communication may operate in a similar way, in that being punished for a particular action may signal disapproval from other group members. Of course, it is not always clear how expectations about preferences translate into actions, particularly where there is a high level of complexity. For example, if one participant is expecting others to behave selfishly, but is unable to discern what the best strategy would be for a self-interested person, then they will have difficulty determining what actions to expect. In cases of confusion, punishment is likely to be less helpful as a signalling mechanism, since those that receive punishment may be no wiser about how to amend their behaviour in accordance with a norm. There is evidence that this also makes participants more hesitant to use punishment. For example, Cason and Gangadharan (2015) have shown that peer punishment is much less frequent in CPR games than in linear VCM games, possibly due to the increased complexity of the game environment making it more difficult establish a social norm. This increased complexity makes it both more difficult to decide whether to punish a participant and more difficult for participants to understand why they are being punished (Cason & Gangadharan, 2015).

Clarity about norms becomes even more difficult when heterogeneity is introduced into the experimental environment. In experiments featuring heterogeneous actors, counter-punishment becomes a much more likely response to enforcement. Counter-punishment occurs when there is normative conflict about what the most appropriate action is for a given decision scenario (Nikiforakis et al., 2012). Reuben and Riedl (2013) tested peer punishment in linear VCM games with various forms of heterogeneity (varying both endowment and MPCR). Using surveys, the authors found that participants generally applied punishment according to one of two heuristics – the "efficiency rule", or the "equality rule". The efficiency rule reflected an expectation that the group should

seek to maximise its returns, regardless of who received payment. The equality rule reflected an expectation that each participant should receive a similar payment, even if this meant that the group was not as productive overall.

In pollution games similar to those of Segerson (1988), Suter et al. (2009) found that a group tax could successfully achieve an efficient pollution level. However, the distribution of abatement between different participants that was observed was sub-optimal. Over- and under-abatement occurred, resulting in a more equal allocation relative to an efficient outcome. Interestingly, these results showed neither the efficiency rule nor the equality rule dominating. It is possible that participants may have found some form of compromise between the two extremes. Suter et al. (2009) proposed that this may have been due to inequality aversion, but did not have survey data to corroborate this. This tendency towards a preference for more equal arrangements than those predicted by efficient levels has also been observed in contracting experiments (Ben-Ner & Putterman, 2009).

Some experiments have explored how endogenous rule formation can be used to counteract some of these challenges and provide clarity about norms and enforcement. Ertan, Page, and Putterman (2009) explored what justifications for punishment might be preferred by allowing participants to formulate rules controlling when punishment was permitted. Even with no prior experience, no group voted to allow unrestricted punishment and no group allowed punishment of high contributors. While groups began with a preference to avoid punishment, over time the punishment of low contributors emerged. Groups with punishment restricted to low contributors exhibited very high levels of cooperation and efficiency (Ertan, Page, & Putterman, 2009; Putterman, Tyran, & Kamei, 2011). These findings align with the literature on communication and covenants, showing the large benefits to cooperation and institutional efficiency when operating on the basis of agreed social norms. More recently, Faillo, Grieco and Zarri (2013) made similar findings in relation to "legitimate punishment", particularly when information is provided about pro-social behaviour from high contributors. Such findings are further supported by experiments showing that players continue to cooperate, even when punishment is not observed, operating

on the assumption and fear that punishment will occur alone (Fudenberg & Pathak, 2010).

Providing further support for these findings, studies of common pool resource management in the field have shown that while homogeneity helps, it is not a necessary requirement for sustaining cooperation. For example, studies from community forestry in Nepal found that heterogeneity was not a strong predictor of the level of collective action. In fact, heterogeneity simply constituted another problem that can be solved with effective institutional design (Varughese & Ostrom, 2001). Collective action can be supported even in highly complex environments, provided that heterogeneous resource users can agree on an appropriate set of rules (Vedeld, 2000).

2.9 Covenants and communication

Communication is being actively omitted from this study, in order to focus more clearly on the impact of sanctions and incentives. This omission necessitates some discussion of the impact of communication on cooperation, in order to clarify what effects might be expected in its absence, particularly in a complex common pool resource management environment.

Communication has long been known to enhance cooperation. Dawes, McTavish and Shaklee (1977) found that even without any kind of sanctioning institution, cooperation was significantly enhanced by relevant communication. Interestingly, when communication was not pertinent to contributions, it had little effect (Dawes et al., 1977). Such findings suggest that that the welfare-enhancing effects of communication are not due merely to group members knowing one another better, but that the use of communication to explicitly develop agreed actions is important. The usual use of communication is in the establishment of agreed contribution (or extraction) levels that support maximisation of returns to all participants (sometimes referred to as “covenants”). Ostrom, Walker and Gardner (1992) showed that such covenants can be effective both with and without an associated punishment. That is, even without the threat of a punishment for

defection, communication and the establishment of an agreement can be a powerful tool for promoting cooperation (Walker, Gardner, Herr, & Ostrom, 2000). In subsequent empirical studies of CPR management groups in the field, communication has been identified as one of the key features of successful groups (Cox, Arnold, & Villamayor-Tomás, 2010).

Meta-analysis of experiments using communication has found a strong positive correlation between communication and cooperation. This is particularly evident in large groups, where cooperation would ordinarily be much more difficult to sustain. While this relationship holds across studies, it is generally strengthened when communication is face-to-face rather than written (Balliet, 2010), though there are exceptions. Bochet, Page and Putterman (2006), for example, found that the use of an anonymous chatroom was almost as efficient as face-to-face communication.

In games exploring communication in the context of contract theory, communication has been found to increase both trust and trustworthiness among participants, as well as overall payoffs. While theory would indicate a preference for binding contracts, experimental findings indicate that with opportunities for communication, contracts in fact become less necessary (Ben-Ner & Putterman, 2009). In studies that combine communication with the use of sanctions, threats are a common use of communication channels. Threats can partially substitute for costly punishment, resulting in higher overall institutional efficiency (Bochet & Putterman, 2009; Vossler et al., 2006). The use of threats extends beyond merely announcing and agreeing contributions. In an experiment by Dugar (2013), the use of approval and disapproval points was substituted for punishment and was found to be effective in increasing contributions illustrating the degree to which social pressure can have a similar effect to monetary sanctions in incentivising pro-social behaviour (Dugar, 2013). However, in contexts where the use of sanctions is available, the credibility of a threat becomes important. Where threats are not followed by sanctions, threats become hollow and ineffective (Masclet, Noussair, & Villeval, 2013).

Cason and Gangadharan (2016) found that in the absence of communication, punishment became relatively ineffective as a mechanism for sustaining cooperation in CPR games, in line with the findings of Ostrom, Walker and Gardner (1992). This is significant, as it goes against findings in linear VCM games (such as those of Fehr and Gächter (2000)) where punishment alone was effective in sustaining cooperation. The implications for institutional design are noteworthy, emphasising the importance of communication and agreed contribution levels in more complex environments. Possible explanations for this finding may be the increased complexity of CPR games that makes optimisation more difficult (confusion) and inadequate punishment ratios limiting the effectiveness of punishment (scofflaws). Further experiments that systematically examine the relationship between incentives and punishment strength are needed to untangle the cause of the observed pattern.

2.10 Critical research gaps of interest

Much is known about the influence of norms, communication and monitoring on levels of cooperation in both public goods and common pool resource games. However, the findings in relation to the use of peer punishment are more mixed between different types of game. In more complex common pool resource environments, it has been suggested that peer punishment is less effective. In the absence of communication, it is possible that the complexity of the environment leads to confusion about norms and the reasons for punishment. In this environment of uncertainty, participants may become confused about the right way to behave in relation to others and may be slow to learn when confronted with punishment. However, these findings are frequently confounded by the lack of a systematic comparison between the cost of punishment and the influence of the payoff environment. The work of Kingsley (2015) suggests that when the levels of complexity and payoffs are consistent across environments, peer punishment has a similar impact on cooperation for both public goods and common pool resources. Given this similarity, it is entirely possible that the use of peer punishment responds to cost and incentive environments in common pool resource games in the same way as has been shown through systemic testing in

VCM environments. As shown by Nikiforakis (2008), reducing the cost of punishment (by increasing the punishment effectiveness ratio) has significant effects on the demand for punishment, the overall levels of cooperation and the degree of welfare that is achieved.

It remains to be seen whether the use of peer punishment in common pool resource games will respond in a similar way. This is a particularly interesting question, as the implication is that the failures of peer punishment that have been seen in common pool resource experiments are in fact an artefact of the different incentives to cooperate or defect, compared to experiments with the VCM. If an exploration of the relationship between incentives and the cost of punishment in a common pool resource environment reveals a similar pattern to that seen in VCM games, there is a positive implication for the potential to sustain cooperation in common pool resource management.

3. Method

The method for this study is set out in a manner typical for economic experiments. We begin by outlining the research gap we are addressing and a central hypothesis. This is followed by a description of the contextual factors around where the experiments were undertaken and the selection of participants. We then describe the economics of our experimental environment, how our sanctioning institutions work and the treatments that have been selected. Finally, we calculate the game theoretic provisions for a cooperative equilibrium (the social optimum) and a non-cooperative (Nash) equilibrium.

3.1 Hypothesis

A number of experiments with non-linear incentive environments have emerged in recent years, featuring peer punishment, group taxes and communication as methods of promoting cooperation. The results have indicated that in the absence of communication, peer punishment is only weakly effective in sustaining cooperation in non-linear environments (Cason & Gangadharan, 2015). It has been hypothesised by these authors that the increased complexity of non-linear environments may make it difficult to identify free riders, and to establish a social norm on which to base punishment decisions. However, it is also important to note that the CPR environment in particular includes much higher incentives to defect, compared to the typical linear VCM commonly used in experiments featuring peer punishment (Cason & Gangadharan, 2015; Fehr & Gächter, 2002). This divergence may change the relationship between punishment effectiveness and participants' willingness to use it. As the ratio of cost to punishment of 1:3 that is typically used was demonstrated as efficient in a linear VCM environment with lower incentives to defect (Nikiforakis & Normann, 2008), a stronger punishment ratio may be required to achieve cooperation in non-linear environments with stronger incentives.

Hypothesis

Increasing the punishment effectiveness ratio in common-pool resource games will result in lower extraction.

This study replicates the recent work of Cason and Gangadharan (2015) featuring peer punishment in a CPR environment, which is then extended by varying the strength of punishment relative to its cost, as well as the incentive environment. We follow the approach of varying punishment ratios used with linear VCM games by Nikiforakis and Normann (2008), in order to test whether more effective punishment can improve cooperation in a non-linear CPR environment. Increasing the strength of punishment is used to determine whether the increased incentives to defect in this design (compared to typical linear VCM games) can be counteracted and increase the overall level of cooperation. We further corroborate our results by testing strong punishment in a higher-payoff environment, to determine whether the change in incentives in a similar non-linear CPR environment has the same effect on cooperation.

3.1 Location and participants

This study was conducted at the Waikato Experimental Economics Lab, at the University of Waikato. We replicate the conditions in Cason and Gangadharan (2015). Individual observations are made of groups of four students, who participate in 10 repetitions of each stage of a CPR game programmed in z-Tree (Fischbacher, 2007). Participants were paid an average rate of $19.34 NZD per hour according to their performance, as is typical for public goods experiments employing university students. Five treatments were conducted, using three different degrees of punishment effectiveness and two different payoff environments (low and high). 11 to 13 independent groups of four are used for each treatment. In total, we have observations from 232 participants, recruited via ORSEE (Greiner, 2004).

3.2 Economic setting

In each period of the game, participants allocate 12 tokens between a private account and a public account. The private return is linear, representing effort expended on the next best alternative to resource use. The payoff from the group account is determined by a quadratic function, representing extraction of a resource with a maximum sustainable yield beyond which group benefits decline. Allocations to the group account impact negatively on other participants, as they represent effort spent extracting a resource that generates declining benefits beyond a maximum sustainable yield. In the context of diffuse losses of nitrate, this would represent the assimilative capacity of a water body. This makes the experiment rival (Apesteguia & Maier-Rigaud, 2006). Due to the concave shape of this group function, both the individual and social optimums are less than the maximum level of extraction, but greater than the minimum.

We use the same design as Cason and Gangadharan (2015), consistent with the CPR literature following Ostrom, Walker and Gardner (1992). Individual payoffs (π) are determined by equation 5:

$$\boldsymbol{\pi_i = C_i + \alpha(E_i - g_i) + \frac{g_i}{G}F(G),} \qquad \textbf{(5)}$$

Where *C* represents a constant (-12), α represents the return from the private account (2), *E* is the initial endowment of 12 tokens and *g* represents the individual's chosen contribution to the group account. The total group contributions $\sum_{j=1}^{n} g_j$ are represented as G. Individuals also receive a share of the group earnings according to their contribution, expressed as $\frac{g_i}{G}$. This makes the group payoff rival, meaning that an individual's consumption of the good reduces others return from consuming the good. In addition, individual free riders are able to capture a large share of the group payoffs. This significantly increases the incentive for defection compared to linear VCM games, which are not rival. In addition to the CPR environment used in Cason and Gangadharan (2015), we also employ a high payoff environment. This enables us to determine the impact of increased incentives to defect on strong punishment, within an otherwise identical

economic environment. The total group payoff *F(G)* is determined by a quadratic function $18G-0.4G^2$ for the low payoff treatments and $24.4G-0.4G^2$ for the high payoff treatments, as shown in figure 2.

Figure 2: The group benefit functions used in the CPR experiments. The maximum of each curve represents the maximum yield from resource extraction.

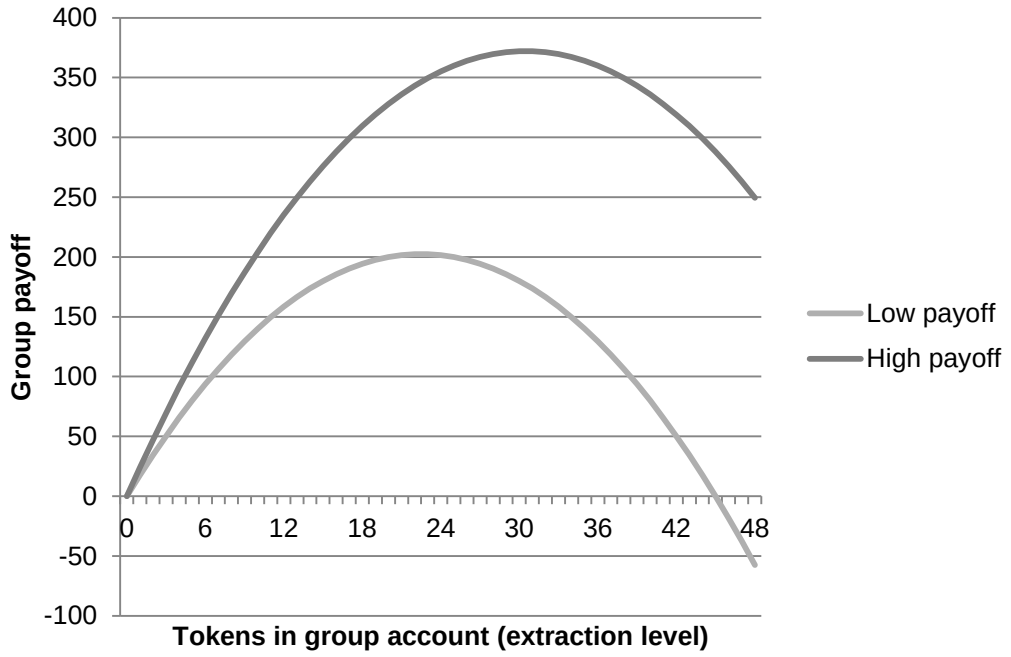

Within these parameters, it is theoretically possible for participants to have negative earnings, particularly after punishment in stage two and particularly for the treatments with higher punishment levels (as for Cason and Gangadharan (2015)). However, none of our 232 participants had negative total earnings in practice. It is straightforward to introduce a constraint that prevents stage one earnings being set to less than zero, as a safeguard against loss of control in this type of experiment. However, this abstracts the experiment further from the reality of the field, where negative earnings for any single year are a relatively common occurrence due to commodity price volatility (Huchet-Bourdon, 2011). Addition of such a constraint would be useful for exploring extreme experimental parameters in future research. For this study, we avoid altering the design to introduce a constraint to ensure comparability with the work we are extending.

Total payoffs for the CPR environement are the combined total of the individual payoffs, plus the total benefits from exploitation of the group resource. Punishment is costly and is experienced as a deadweight loss to the group as a whole, as is typical in the literature. The cost is the sum of the assigned punishment points (*Pa)* and the punishment points received (*Pr*) multiplied by the punishment ratio *R*. The overall net social benefits accordingly include both the gains from resource use and the cost of punishment, as shown in equation 6. We employ identical punishment institutions across the treatments, varying only the strength of the punishment relative to its cost. Table 3 summarises the features of the economic environment.

The overall net social benefits (*NSB*) for the CPR environment can therefore be expressed as:

$$NSB = \sum_{i=1}^{n_i} -12 + 2(12 - gi) + \left[18\sum_{i=1}^{n_i} g_i - 0.4\left(\sum_{i=1}^{n_i} g_i\right)^2\right] - \left[\sum_{i=1}^{n_i} Pa_i + RxPr_i\right] \quad (6)$$

Table 3: Outline of key features of the experimental environments

	Low Payoff	**High Payoff**
Private return	2 per token	2 per token
Group return	$(g_i/G)(18G-0.4G^2)$	$(g_i/G)(24.4G-0.4G^2)$
Punishment ratio (*R*)	None, 1:3, 1:5, 1:9	1:9
Constant	-12	-12
Nash equilibrium (payoff)	8 (37.6)	11 (64.8)
Social optimum (payoff)	5 (52)	7 (90.4)
% Gain from cooperation	38.3%	39.5%
Maximum payoff from defection	22.4	32
Observations per treatment	11/44 to 13/52 per treatment	11/44

3.4 Punishment

Peer punishment is a feature of most of the treatments undertaken. This means that players in each group can fine other members of their group after seeing the extraction levels of all group members. We employ the approach that is commonly used in linear VCM environments. This allows all participants to assign punishment points to any and all players, with no limitations (such as only being able to punish defectors). Participants can assign up to five punishment points to each of the other members of their group, at a cost of 1 experimental dollar per point. Initially, the cost-to-punishment ratio is 1:3, meaning a fine of 3 experimental dollars per point. This means that any individual can assign a fine of up to 15 experimental dollars to each other individual in their group. This ratio of 1:3 has been demonstrated to be the minimum level required to halt the decline in cooperation over time in linear VCM experiments and has been in common usage subsequently (Cason & Gangadharan, 2015; Nikiforakis & Normann, 2008).

The incentives to defect are very different in the CPR environment compared to the linear VCM environment. Not only are the incentives to defect much higher, they are non-linear (see comparison in figure 3) with stronger or weaker incentives to defect at different extraction levels. In addition, the amount of damage experienced by cooperative group members due to the defection of one member is more severe. This changes the ratio of gain received (by the defector) to damages caused (to cooperators) by a defection in the CPR game compared to the VCM. These ratios are compared in figure 4.

These differences in the incentives to defect and the consequences of defection for others may be affecting the relationship between the cost of punishment and the level of extraction that occurs. In order to provide a logic for the selection of different punishment ratios, we examine two different explanations of what is necessary for punishment to be effective. This enables us to cover a much greater range in punishment effectiveness across our treatments, while not needing to undertake experiments for every possible level of punishment in between. The logic of our explanations focuses on key differences in the gains to defectors and

damage to cooperators in the two different incentive environments of the CPR and the linear VCM game. To choose each ratio, we examine the incentive structure in the linear VCM and compare it to the incentive structure in the CPR. We then multiply the punishment effectiveness to a point where it is as strong relative to the incentives in the CPR as it would be in the linear VCM if it were set at 1:3. Because of the complex non-linear relationships in incentives between the two games, there are multiple explanations for how the ratio of 1:3 might be interacting with particular features of the linear VCM environment to produce results where cooperation does not decline over time. These explanations are expanded below and translated into levels of punishment effectiveness for use in our CPR games.

Figure 3: Comparison of payoffs from defection in linear VCM and low payoff CPR games[2]

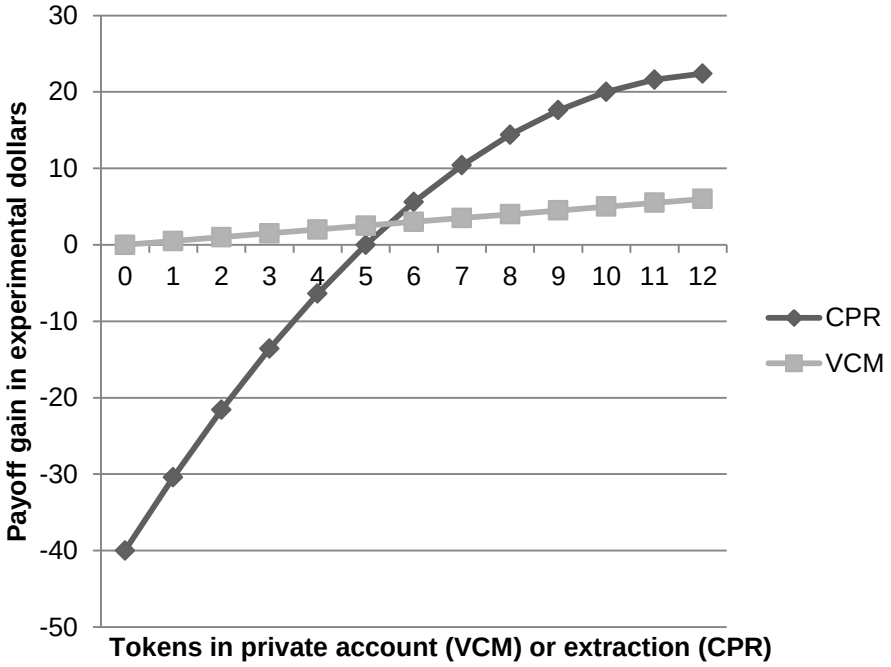

[2] This assumes the defection of a single individual in relation to three other cooperative individuals playing strategies consistent with the social optimum. This wouls mean zero tokens in the private account for the linear VCM and 5 tokens of extraction for the CPR game (evident from the x intercepts of the two curves).

Figure 4: Comparison of ratios of gain to damage caused by defection in linear VCM and low payoff CPR environments[3]

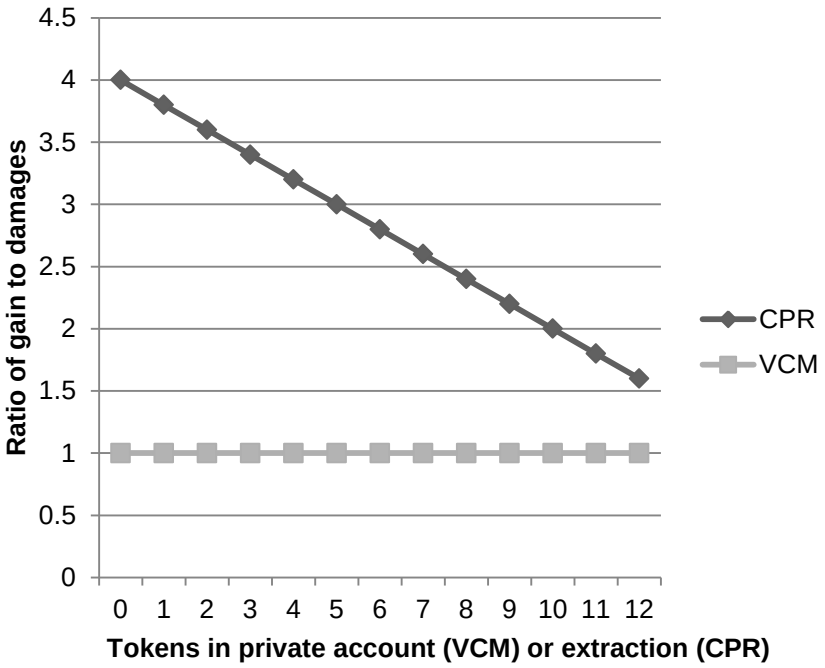

Explanation 1: Punishment levels respond to damage levels

Our first possible explanation for the difference between the VCM and the CPR is that cooperators' willingness to use punishment is related to the amount of damage that they receive, relative to the gains of defectors. This kind of punishment behaviour would indicate a greater willingness to use punishment, in order to counteract greater damage caused by defection. There are several ways in which this might play out in relation to the different incentives between the linear VCM and the CPR.

First, we examine the ratio of gains to damages for full defection (a contribution of 12 in the CPR environment and 0 in the linear VCM). This gain to damage ratio is 1:1 in the linear VCM and 1:1.6 in the CPR environment. Given that a punishment effectiveness ratio three times this is required to halt the decline in cooperation in the linear game, we apply the same muliplier to the incentive in the CPR environment. In this case, this implies a punishment effectiveness ratio of

[3] As for figure 3, this assumes the defection of a single individual in relation to three other cooperative individuals playing strategies consistent with the social optimum. This would mean zero tokens in the private account for the linear VCM and 5 tokens of extraction for the CPR game.

1:4.8 to maintain the same relative strength to the linear VCM. We round this up to 5 to simplify the parameters of the experiment for participants. In our design, this punishment effectiveness ratio also reflects the difference in the average ratio of marginal gain to marginal damages between the linear VCM and CPR. The design allows us to test both possible requirements for punishment effectiveness simultaneously.

Second, we compare the maximum ratio of marginal gain to marginal damages that can occur in the two games. At 1:2.8, the ratio is significantly higher in the CPR game, yielding a punishment ratio of 9 (rounded up from 8.4 to ensure sufficient strength and simplicity). These three approaches to defining punishment ratios are summarised in table 4.

Explanation 2: Prospering cheaters

Our second possible explanation is that willingness to use punishment is not directly related to the level of damages received when a group member defects, and that the likelihood of defection paying off is the primary driver of punishment effectiveness. This would mean other factors besides damages drive the amount of punishment players are willing to use, and whether this level of punishment is effective or not is related to the impact that it has on the earnings of defectors. In order to tease this apart, our high payoff CPR treatment is designed with similar marginal incentives to the low payoff treatments, but different absolute benefits from defection. For example, if a punishment ratio of 1:9 is successful in the low payoff environment, this could mean either marginal gain-to-damage incentives or prospering cheaters drive punishment effectiveness. Success of this ratio in the high payoff environment would indicate that marginal gain-to-damage relationships are the most important driver of punishment effectiveness. However, a failure of this ratio in the high payoff environment would indicate that the dominant driver is simply the degree of benefits from defection (as shown in table 4).

Table 4: Differences in incentives between linear VCM and CPR environment and associated punishment ratios

Normal payoff CPR

	Linear VCM	CPR	CPR relative to VCM	Punishment ratio
Max gain:damage ratio	1	1.6	4.8	5
Max marginal gain:marginal damage ratio	1	2.8	8.4	9
Highest possible payoff through defection	50	146.4	8.78	9

High payoff CPR

	Linear VCM	CPR	CPR relative to VCM	Punishment ratio
Max gain:damage ratio	1	2.86	8.57	9
Max marginal gain:marginal damage ratio	1	2.86	8.57	9
Highest possible payoff through defection	50	223.20	13.39	14

3.5 Treatments

For this study, the ratio of punishment cost to punishment effectiveness is used as a treatment variable. This approach follows Nikiforakis and Normann (2008), who found that a punishment ratio of 1:3 was most efficient in linear VCM games. Cason and Gangadharan (2015) employed the same 1:3 ratio and found it to be weakly effective in the CPR environment. For this study, 1:3, 1:5, and 1:9 are used as alternative ratios, in order to explore whether more effective punishment can counter the increased incentives to defect in non-linear environments. After each round, participants can see all decisions made by all participants. They then have the opportunity to punish any other participants with

0-5 punishment points. In order to further differentiate between our two hypotheses, a fifth treatment is undertaken in a higher incentive environment, with similar marginal incentives to defect. The five treatments are summarised in table 5.

Table 5: Treatments

Treatment	Description
No punishment [CPRP0]	Control treatment of CPR game without opportunities for punishment (replication of Cason and Gangadharan 2015).
1:3 [CPRP3]	Peer punishment with 1:3 punishment effectiveness ratio (replication of Cason and Gangadharan 2015). Target any players with 0-5 punishment points. Cost to punish is 1 experimental dollar per point and imposes a fine of 3 experimental dollars per point.
1:5 [CPRP5]	Peer punishment with 1:5 punishment effectiveness ratio. Target any players with 0-5 punishment points. Cost to punish is 1 experimental dollar per point and imposes a fine of 5 experimental dollars per point.
1:9 [CPRP9]	Peer punishment with 1:9 punishment ratio. Target any players with 0-5 punishment points. Cost to punish is 1 experimental dollar per point and imposes a fine of 9 experimental dollars per point.
1:9 with high payoff [CPRP9H]	Higher-payoff from resource extraction. Peer punishment with 1:9 punishment effectiveness ratio. Target any players with 0-5 punishment points. Cost to punish is 1 experimental dollar per point and imposes a fine of 9 experimental dollars per point.

3.6 Game Theoretic Predictions

Low payoff environment

In the absence of cooperation, we expect individuals to maximise their own benefits, assuming other participants will do the same, with the system converging to the Nash equilibrium. For our CPR game design, there are strong incentives to defect, with a payoff of 22.4 experimental dollars for defection from the optimum

in the lower payoff environment. The individual payoff function from our low payoff CPR economic setting is shown in equation 7:

$$\pi_i = -12 + 2(12 - g_i) + \frac{g_i}{\sum_{j=1}^{n} g_j}\left[18\sum_{j=1}^{n} g_j - 0.4\left(\sum_{j=1}^{n} g_j\right)^2\right] - Pa - RxPr \quad (7)$$

Analysing the problem with the assumption of rational actors, we can assume punishment will not occur. This is because punishing other individuals comes at a cost to the punisher, while conferring no monetary benefit. Under these conditions, rational individuals are not predicted to use punishment. This allows us to ignore this part of the equation in calculating the Nash choice. As the function is strictly concave, a maximum occurs at $\pi_i' = 0$, which in this case provides a value for the Nash choice of 8 tokens contributed to the public account (extraction from the CPR). The Nash equilibrium provides net social benefits of 150.4, as shown in table 6. In a CPR game, the Nash choice represents a state where the individual motivations of participants drive them towards over-exploitation of the common resource. This means that resource extraction has reached a level beyond peak returns, generating sub-optimal returns to each individual.

A social (Pareto) optimum occurs when all participants cooperate in order to maximise total social benefits. This involves optimising all levels of emissions with the sole target of maximising net total benefits. By cooperating and reducing overall extraction of the CPR, participants are able to achieve net social benefits of 208 if choosing the social optimum amount of 5 tokens, as shown in table 7. Returns to individuals are maximised, as the resource extraction is generating the maximum sustainable yield. However, it is important to note that if the other three group members make choices at this social optimum, an individual defector is able to capture net benefits of 74.4 experimental dollars, giving gains for defection of 22.4.

Table 6: Nash equilibrium predictions for low payoff games

	Ei	*gi*	Individual benefits	Group benefits	Net benefits
Individual	12	8	-4	41.6	37.6
TOTAL	48	32	-16	166.4	150.4

Table 7: Social optimum predictions for low payoff games

	Ei	*gi*	Individual benefits	Group benefits	Net benefits
Individual	12	5	2	50	52
TOTAL	48	20	8	200	208

High payoff environment

In our high payoff environment, there are even stronger incentives to defect, with a payoff of 32 experimental dollars for defection from the optimum. The individual payoff function from our CPR economic setting is shown in equation 8:

$$\pi_i = -12 + 2(12 - g_i) + \frac{g_i}{\sum_{j=1}^{n} g_j}\left[24.4\sum_{j=1}^{n} g_j - 0.4\left(\sum_{j=1}^{n} g_j\right)^2\right] - Pa - RxPr \quad (8)$$

As for the low payoff treatments, we can assume punishment will not occur if the assumption of rationality holds. This assumption allows us to ignore this part of the equation in calculating the Nash choice and solve for a maximum at $\pi_i' = 0$, which in this case provides a value for the Nash choice of 11 tokens contributed to the public account (extraction from the CPR). The Nash equilibrium provides net social benefits of 259.2, as shown in table 8.

In the high payoff treatment, participants are able to achieve net social benefits of 361.6 if choosing the social optimum amount of 7 tokens, as shown in table 9. In the case of a single defector from the optimum in this high payoff environment, an individual is able to capture net benefits of 122.4 experimental dollars, giving gains for defection of 32.

Table 8: Nash equilibrium predictions for high payoff games

	Ei	***gi***	**Individual benefits**	**Group benefits**	**Net benefits**
Individual	12	11	-10	74.8	64.8
TOTAL	48	44	-40	299.2	259.2

Table 9: Social optimum predictions for high payoff games

	Ei	***gi***	**Individual benefits**	**Group benefits**	**Net benefits**
Individual	12	7	-2	92.4	90.4
TOTAL	48	28	-8	369.6	361.6

4. Results

Results are structured according to observations in relation to extraction levels (contributions to the group account), the use of punishment, how participants respond to receiving punishment and efficiency. A brief discussion of some of the strategies employed by individual participants is also included to provide a sense of how the high-level data is formed from the combination of individual strategies.

4.1 Extraction levels

If the cost and effectiveness of punishment does not affect the overall levels of extraction observed in CPR games, we would expect to find no difference between treatments, provided the incentives for cooperation and defection remained the same across economic environments. In our treatments, marked differences were found in the average levels of extraction at each different punishment effectiveness ratio. Differences also occurred in the degree of extraction over time. Higher punishment effectiveness ratios produced much more constant levels of extraction over time.

Result one: Stronger punishment results in lower extraction levels

We find significantly lower extraction levels associated with higher punishment ratios. This decrease in extraction is modest but statistically significant and monotonic. All punishment treatments had significantly lower extraction levels than those with no punishment, with stronger punishment showing a more significant difference (first row of table 11). Though a ratio of 1:5 was not significantly different to a ratio of 1:3 ($p=0.447$), a ratio of 1:9 achieved lower contributions compared to all other treatments (as shown in the right-hand column of table 11). Figure 5 shows the extraction results over time. Note the stacked lines created by the increasing effectiveness of punishment in reducing extraction from the low payoff CPR environment. Despite the significant improvement with stronger punishment, extraction levels were still considerably higher than the social optimum of 5, with the lowest mean contribution achieved being 6.86 with

a punishment ratio of 1:9 (table 10) in the low payoff environment. Behaviour was less responsive to increased punishment strength compared to results in the linear VCM (Nikiforakis & Normann, 2008).

Table 10: Mean extraction levels

Standard errors are in parentheses. P-values for Wilcoxon signed-rank tests of difference in public contribution levels between the first and last five periods are in the right column.

Treatment	Periods 1-10	Periods 1-5	Periods 6-10	Wilcoxon signed-rank p-value
No punishment	7.59 (2.60)	7.45 (2.78)	7.73 (2.42)	0.706
1:3 punishment	7.26 (2.60)	7.13 (2.64)	7.40 (2.57)	0.097
1:5 punishment	7.19 (2.46)	7.10 (2.56)	7.28 (2.35)	0.637
1:9 punishment	6.86 (2.56)	6.86 (2.60)	6.87 (2.51)	0.549
1:9 punishment high payoffs	8.81 (2.72)	8.29 (2.74)	9.33 (2.61)	0.001

Table 11: Differences in extraction between treatments

P-values are for Mann-Whitney tests (2-tailed) of difference in extraction levels between punishment treatments.

	1:3 punishment	1:5 punishment	1:9 punishment
No punishment	0.052	0.006	<0.001
1:3 punishment	.	0.447	0.001
1:5 punishment	.	.	0.011

Figure 5: Mean extraction - low payoff treatments

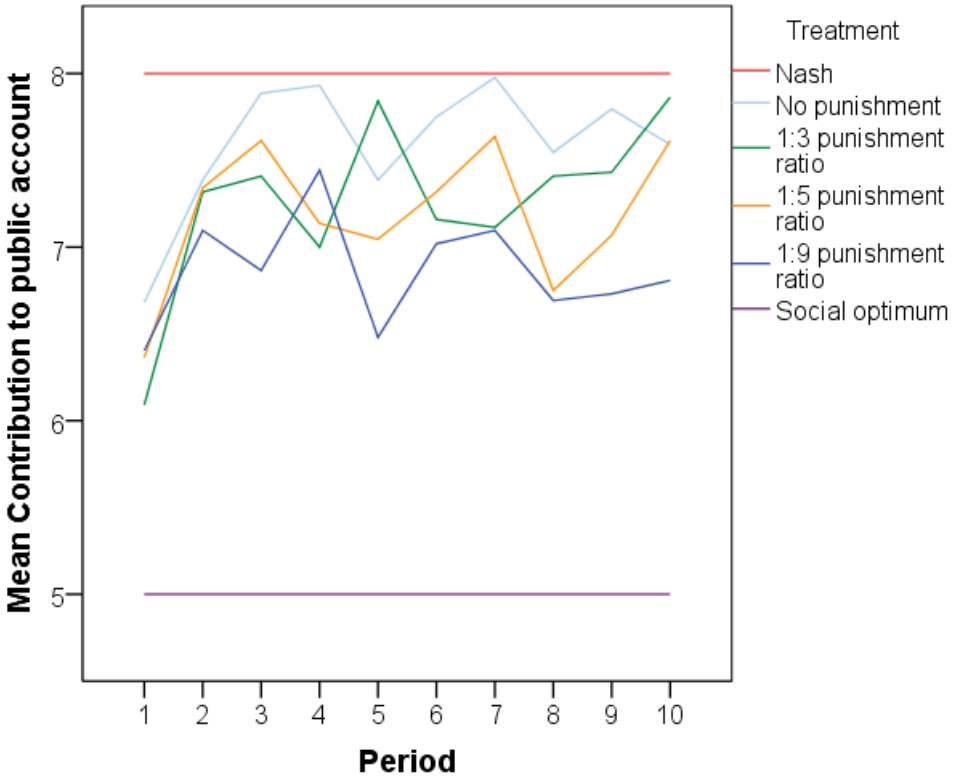

Figure 6: Mean extraction comparing high and low payoff environments

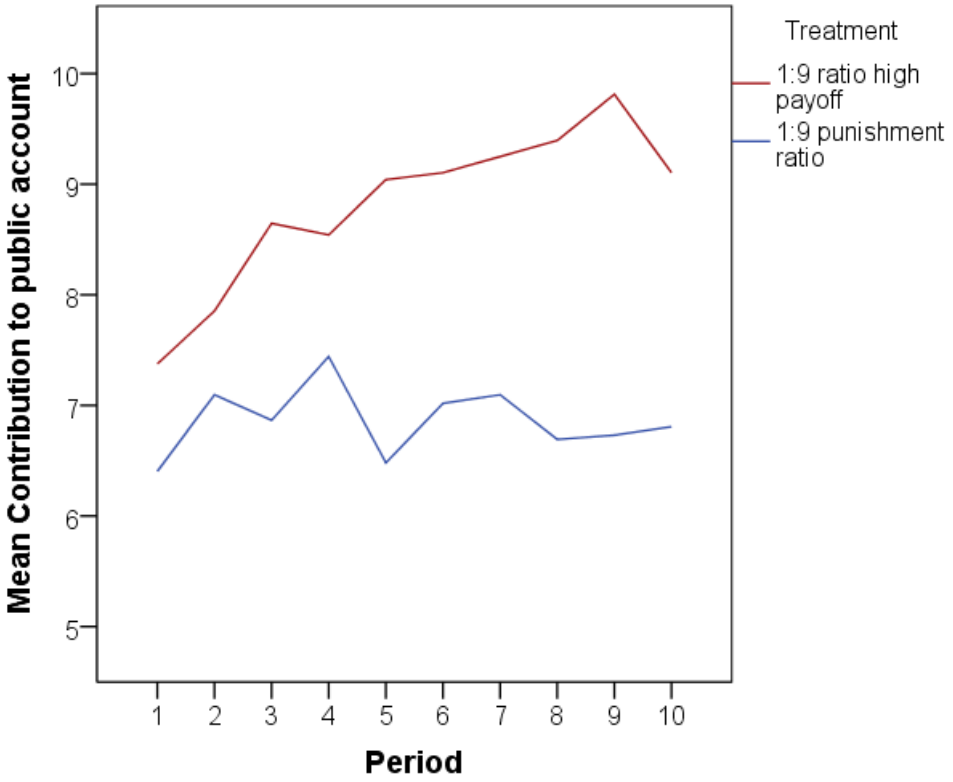

Result two: Stronger punishment inhibits decline in cooperation over time

Although average extraction remained lower in the treatments with punishment than those with no punishment, several of our treatments were characterised by a decline in cooperation over time, with higher extraction in the later periods. This outcome was particularly marked in the high payoff treatment (see tables 10 and 12 and figures 6 and 7). Spearman correlations in table 12 indicate significant declines for both the 1:3 punishment treatments and the 1:9 high payoff treatments.

Declines were not significant for the no punishment, 1:5 and 1:9 ratio treatments. In the case of the no punishment treatment, this is likely to be due to the fact that extraction quickly reached a higher level than in the other treatments, remaining relatively flat but close to the Nash equilibrium. In the case of the 1:5 ratio treatment, decline appears to be somewhat curbed by punishment. Extraction in the 1:9 punishment treatment was particularly flat with a Spearman correlation coefficient of -0.003 and a p-value of 0.952, as shown in table 12. Table 10 also compares the mean extraction levels from the first and last five periods, including Wilcoxon signed-rank tests for significance of differences between early and late periods. Figure 6 illustrates this information, showing the especially marked decline in the higher payoff environment.

Table 12: Correlations in extraction over time

P-values for Spearman correlations (2-tailed) across all 10 periods

Treatment	Spearman p-value
No punishment	0.215
1:3 punishment ratio	0.008
1:5 punishment ratio	0.319
1:9 punishment ratio	0.952
1:9 punishment ratio high payoffs	<0.001

Figure 7: Mean contributions - low payoff treatments

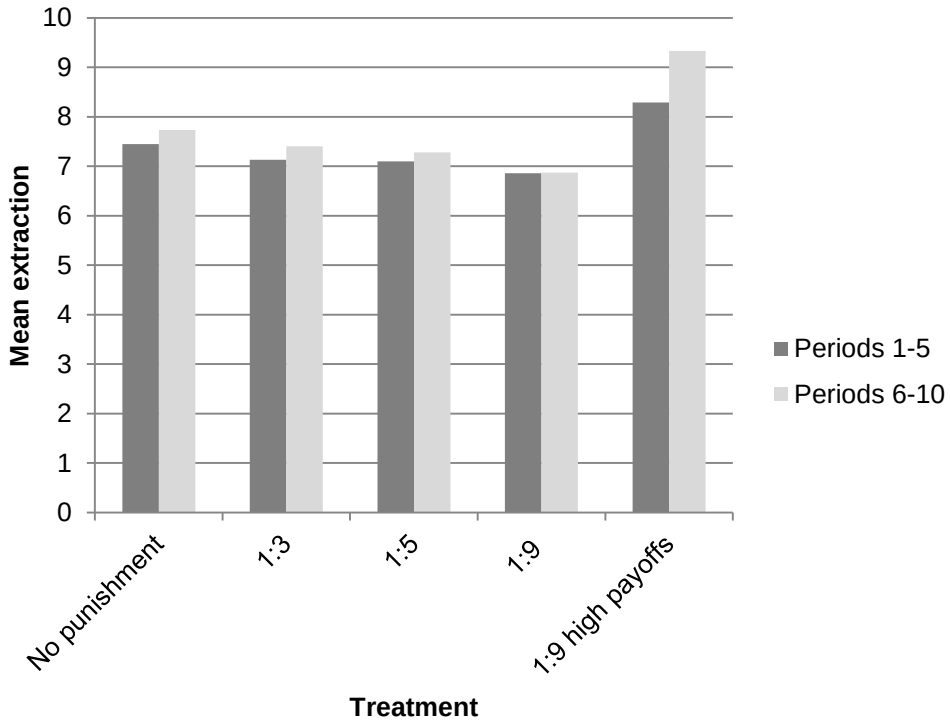

4.3 Punishment – use

Result three: More effective punishment and higher incomes generate greater demand for punishment (punishment is a normal good)

In all cases, higher punishment effectiveness induced greater use of punishment, with significantly greater levels of punishment observed in stronger punishment treatments. This is evident in figure 8, with significance established in table 14. In effect, increasing the punishment ratio reduces the cost of a given amount of deduction, which generates increased demand. This finding is consistent with previous findings in relation to punishment as a normal good; for example, Ostrom et al. (1992) and Nikiforakis & Normann (2008). In addition to the impact of reduced cost, increased income (as observed in the high payoff treatment) had a further effect of increased demand for punishment.

The observed level of punishment over time differed markedly between treatments (tables 13 and 14). For the 1:3 punishment ratio, a relatively low level of punishment occurred and was constant over time. Increasing this ratio to 1:5

produced initially high levels of punishment, but a decline in use over time. Given that this decline in punishment is not associated with a significant increase in extraction for this treatment, the implications for efficiency are positive (and indeed evident in the significant increase in efficiency over time in table 17). The declining level of punishment in the 1:5 treatment contrasts with the relatively constant level of punishment observed in the 1:9 treatment. Efficiency in the 1:5 treatment began lower than in the 1:9 treatment due to higher extraction levels, but was not significantly different in the latter periods (as shown in table 18) due to the decline in punishment costs.

Table 13: Mean payoff reductions received

Standard errors are in parentheses. P-values for 2-tailed Wilcoxon signed-rank tests of difference in between the first and last five periods are in the right hand column.

Treatment	Periods 1-10	Periods 1-5	Periods 6-10	Wilcoxon signed-rank p-value
1:3 punishment	4.40 (5.89)	4.79 (5.70)	4.01 (5.89)	0.229
1:5 punishment	9.81 (11.94)	12.25 (12.63)	7.36 (10.69)	<0.001
1:9 punishment	10.94 (16.93)	10.07 (15.87)	11.80 (17.91)	0.342
1:9 high payoff	17.89 (20.55)	15.56 (17.84)	20.21 (22.75)	0.052

Higher incomes in the high payoff treatment significantly increased the amount of punishment participants were willing to pay for. However, the quantity of punishment demanded increased less than the associated increase in income, in the high payoff treatment. Accordingly, the income elasticity of demand was relatively low at 0.66.

By contrast, demand was highly-responsive to price for the initial drop in cost from a 1:3 to a 1:5 punishment ratio. Price elasticity appears to fall rapidly above this level, however, as shown by the demand curve in figure 9. It seems likely that the quantity of punishment demanded would continue to increase as the

punishment ratio increases further, up to some maximum, but this is an open empirical question left for further research.

Figure 8: Payoff reductions received in each period

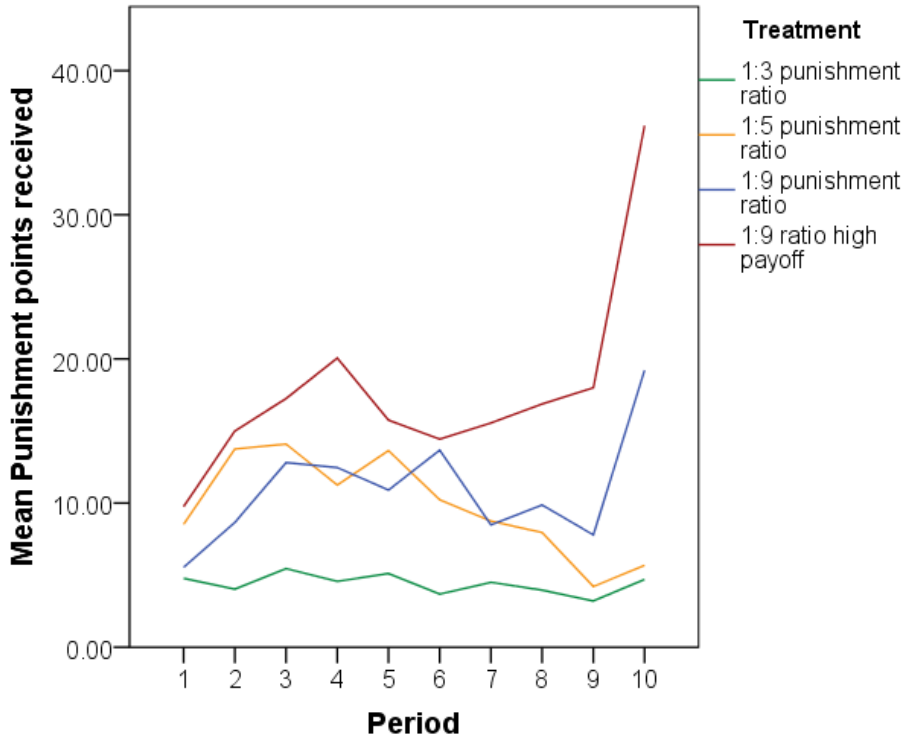

Figure 9: Demand for punishment at 1:3, 1:5 and 1:9 effectiveness ratios

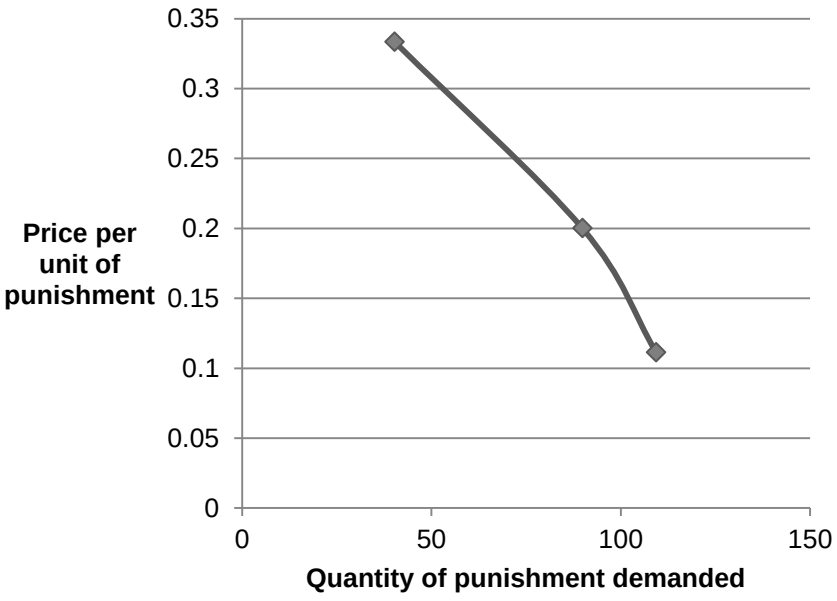

Table 14: Differences in punishment use

P-values for Mann-Whitney tests (2-tailed) of difference in punishment levels between treatments.

Periods 1-10			
	1:5 punishment	1:9 punishment	1:9 punishment high payoff
1:3 punishment	<0.001	<0.001	<0.001
1:5 punishment	.	0.044	<0.001
1:9 punishment	.	.	<0.001
Periods 1-5			
	1:5 punishment	1:9 punishment	1:9 punishment high payoff
1:3 punishment	<0.001	0.111	<0.001
1:5 punishment	.	<0.001	0.519
1:9 punishment	.	.	<0.001
Periods 6-10			
	1:5 punishment	1:9 punishment	1:9 punishment high payoff
1:3 punishment	0.035	<0.001	<0.001
1:5 punishment	.	0.135	<0.001
1:9 punishment		.	<0.001

Result four: More effective punishment is used to increasingly target defectors

The distribution of punishment points that were assigned demonstrated a clear pattern of targeting defectors, as shown in figure 10. Because the social optimum (and average contribution) is at a level of extraction above the minimum but below the maximum, we also observe punishment being assigned to those contributing less than the group average. While it is tempting to attribute this result to group norm enforcement, such a claim cannot be made definitively. For example, a number of participants used strategies of maximum extraction (all 12 tokens) coupled with assigning the maximum number of punishment points to every other participant, in order to enforce asymmetrical extraction levels within

the group. While the strategy is certainly not random, it could not be described as enforcing a shared norm. Further information is required regarding participants' attitudes to norms and punishment to decouple these possible explanations.

Figure 10: Payoff reductions received

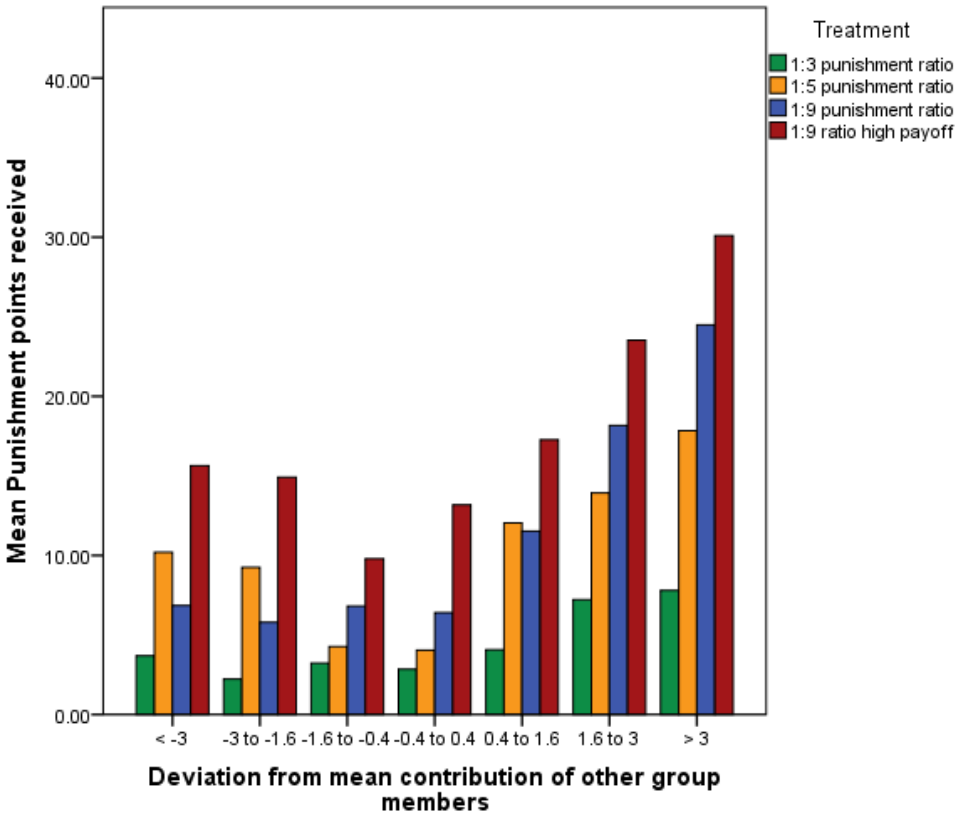

4.4 Punishment – impact

Result five: The direct impacts of strong punishment interact with the indirect effects

While stronger punishment induced lower extraction levels monotonically, the results are less clear cut for the direct relationship between the occurrence of punishment and a participant's response in the following period. Across all treatments, the use of punishment was significantly negatively correlated with changes in extraction in the following period (as shown in table 15 below). This means that, in general, the most likely response to punishment was to reduce extraction in the next period. However, the strongest correlation was in fact in the weakest punishment treatment. This strong correlation may be due to aversion to

the risk of being punished playing a greater part than the direct impact of punishment in the stronger punishment treatments. That is, it may be that weak punishment operated primarily as a social signal, which was then crowded out by more effective sanctions in the 1:5 treatment, with the 1:9 treatment introducing sanctions effective enough to overcome this crowding out effect. Similar patterns have been documented in relation to the use of incentives (Gneezy et al., 2011). Alternatively, the correlation may be an artefact of lower overall extraction in the lower punishment treatments. When group means are lower, the payoff for each unit of defection is higher, due to the non-linear group return function. This may make risk-loving individuals less likely to respond to punishment. Despite the noisy relationship between punishment strength and the response to punishment, the strongest behavioural amendments clearly occurred where the strongest punishment occurred, as shown in figure 12. The overall correlation between payoff reduction and change in extraction was significant at -0.143 (p-value <0.001).

Figure 11: Payoff reductions received and mean reaction

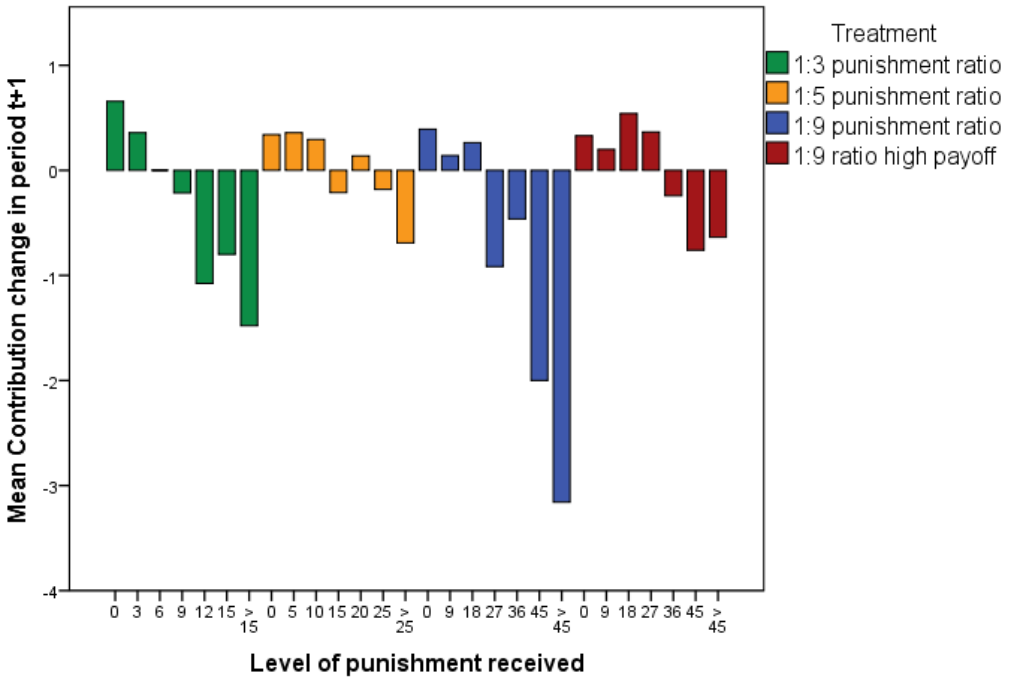

Table 15: Correlations between punishment and change in extraction in the following period

P-values a for partial Spearman correlations (2-tailed), controlling for the effect of period.

Treatment	Spearman correlation (p-value)
1:3 punishment ratio	-0.214 (0.001)
1:5 punishment ratio	-0.103 (0.040)
1:9 punishment ratio	-0.195 (<0.001)
1:9 punishment ratio high payoffs	-0.113 (0.019)
All punishment treatments	-0.143 (<0.001)

Result six: More effective punishment can remove the capability for defectors to earn more than co-operators

Increasing the strength of punishment significantly alters the distribution of earnings in relation to participants' deviations from their group mean. Table 16 shows the results for 2-sample Kolmogorov-Smirnov tests establishing significance at the 95 or 99 per cent level for most treatments. The single exception is the 1:3 and 1:9 treatments, although it should be noted that a Mann-Whitney test returned a p-value of 0.048 for this combination.

In the low payoff environment, the highest earners were also the largest defectors for both the no punishment and 1:3 punishment treatments. This finding indicates that given the demand for punishment at a ratio of 1:3, insufficient punishment occurs to remove the monetary incentive to defect, as the earnings from these defections are big enough to offset the likely amount of punishment that occurs. This outcome was somewhat moderated by increasing the ratio to 1:5, but the highest defectors still earned more than lower defectors. It was only by increasing the punishment ratio to 1:9 that we were able to induce a monotonic decrease in earnings associated with higher levels of defection. This is clearly evident in figure 12. Transferring the 1:9 punishment ratio into our high payoff environment, the relationship between defection and earnings returned, as shown in figure 13.

Figure 12: Earnings from defection - all low payoff treatments

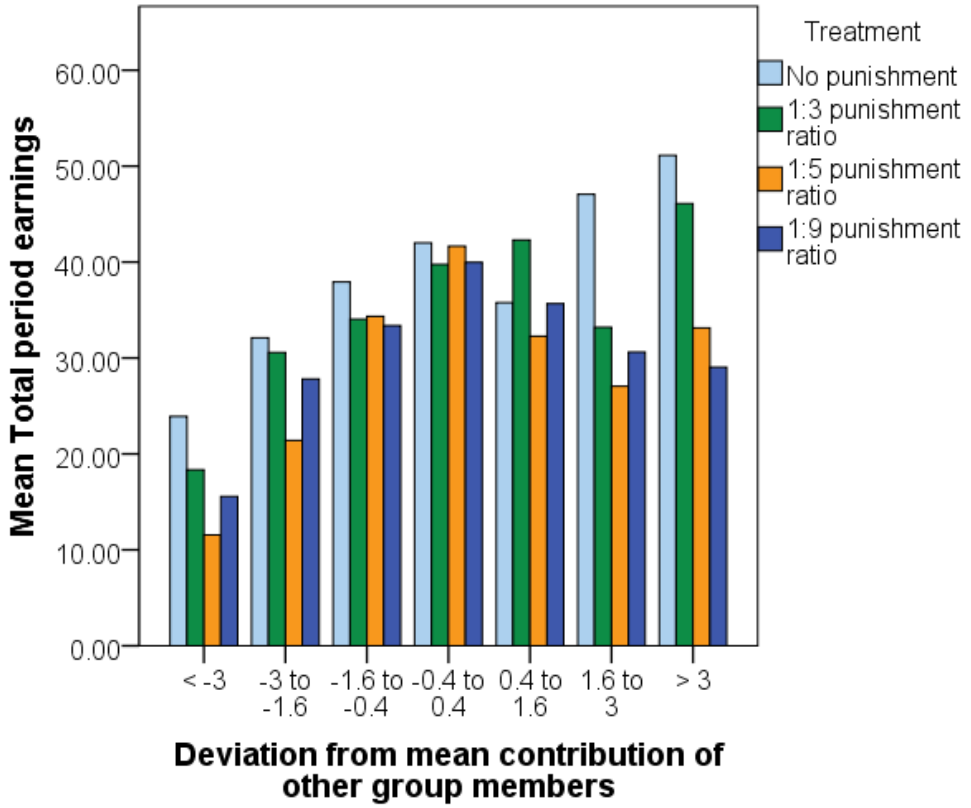

Figure 13: Earnings from defection - comparison with high payoff treatment

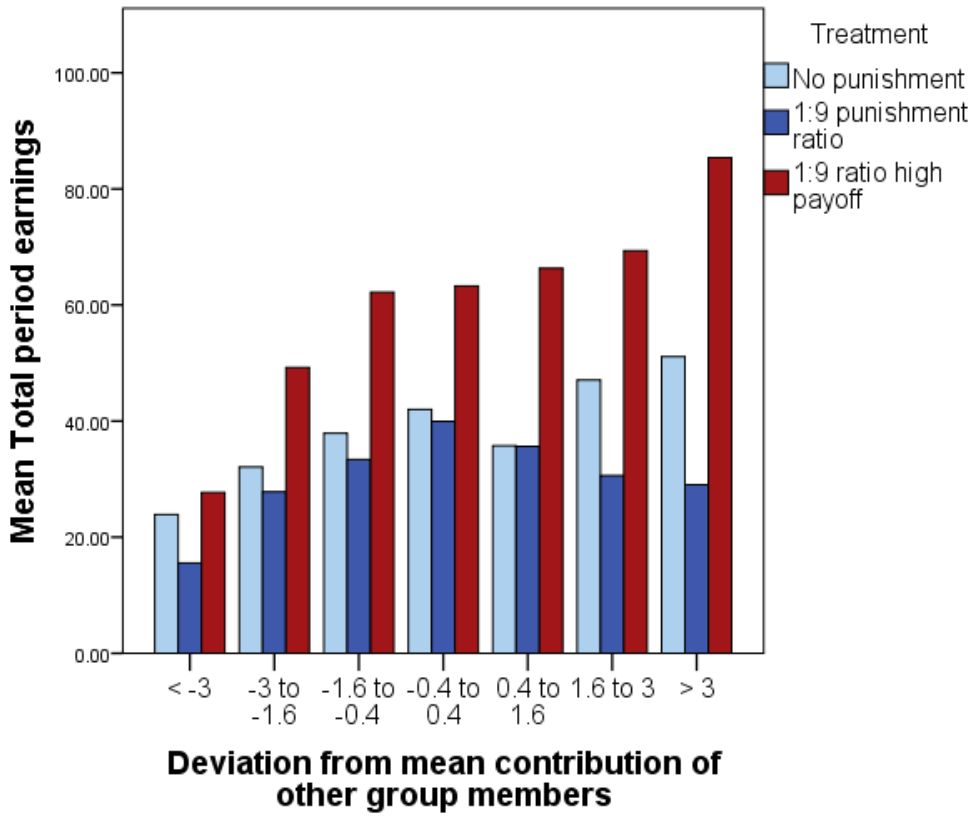

Table 16: Differences in earnings due to punishment

P-values for 2-sample Kolmogorov-Smirnov tests for difference in the distributions of earnings between punishment treatments in the low payoff environment (the high payoff treatment is excluded due to earnings not being directly comparable).

Periods 1-10			
	1:3 punishment	1:5 punishment	1:9 punishment
No punishment	<0.001	<0.001	<0.001
1:3 punishment	.	<0.001	0.061
1:5 punishment	.	.	0.021

4.2 Efficiency

Efficiency refers to the overall welfare gains or losses that occur relative to the self-interested Nash equilibrium. The efficiency calculation provides an overall measure of the welfare impact of a particular institution on an incentive environment. In our case, welfare is a combination of the yields of individuals from their private accounts and the yields from the common pool resource, less the costs of punishment.

Efficiency is calculated as a proportion of the potential gains from achieving the social optimum, relative to the Nash equilibrium earnings, as shown in equation 9.

$$\boldsymbol{Efficency} = \frac{\boldsymbol{\pi_{actual} - \pi_{Nash}}}{\boldsymbol{\pi_{Soc.Op.} - \pi_{Nash}}} \qquad \textbf{(9)}$$

In our design, punishment is costly and does not convey any direct welfare benefit. Because actual earnings are dependent on the level of punishment points assigned and received, efficiency is a function of both extraction and the level of punishment that occurs. Punishment effectiveness, the amount that is used and its impact on decisions combine to influence overall efficiency (Nikiforakis & Normann, 2008).

Result seven: The cost of increasingly severe punishment offsets the welfare gains of lower extraction (at least in the short-term)

All treatments featuring punishment exhibited initial declines in efficiency followed by gradual improvements, with a drastic decline in the last period due to high levels of punishment (as shown in figure 14). This feature is interesting in and of itself, in that no amendment to behaviour can be expected to be derived from punishment in this last period. Final period punishment appears to be the product of participants exercising spite, due to the lack of a fear of reprisal in the last period. As shown in table 17, average efficiencies are negative for all treatments with punishment. There is a significant improvement in efficiency for the 1:5 punishment treatments and a significant decline in the high payoff treatment. These differences remain even when omitting the final period from analysis.

Table 18 outlines the statistics for differences between treatments. As is evident in figure 14, efficiency levels are in fact highest in the no punishment treatment, despite this treatment having the highest overall levels of extraction. All punishment treatments have significantly lower efficiency than the no punishment treatment, with some having significantly lower efficiency compared to other punishment treatments. In order of descending efficiency, results were no punishment, 1:9 punishment with high payoffs, 1:3 punishment, 1:9 punishment and 1:5 punishment. Interestingly, no significant difference in efficiency was observed between punishment treatments for the last five periods. The combination of results suggests that punishment, earnings and overall efficiency interact dynamically. For example, high earnings generate high efficiency initially in the high payoff treatment, but the ineffectiveness of a punishment ratio of 1:9 in this environment leads to a dramatic decline in efficiency as extraction levels increase and earnings decline.

As demonstrated above in result one, stronger punishment is effective in reducing overall extraction levels. Stronger punishment increases the earnings from the resource and the total stage one earnings for participants. However, a significant amount of punishment is used to generate these lower extraction levels. As greater punishment occurs in higher-punishment treatments, the modest gains in

earnings are offset by the increased costs from punishment. This result is in stark contrast to the findings of Nikiforakis and Normann (2008) when undertaking comparative statics analysis of punishment effectiveness in linear VCM games. Nikiforakis and Normann (2008) found clear, monotonic improvements in efficiency associated with increasing punishment effectiveness (particularly for latter periods). The increased complexity of the CPR environment in our design generated no improvements in efficiency compared to our control treatment, even in latter periods.

Figure 14: Mean efficiency relative to social optimum

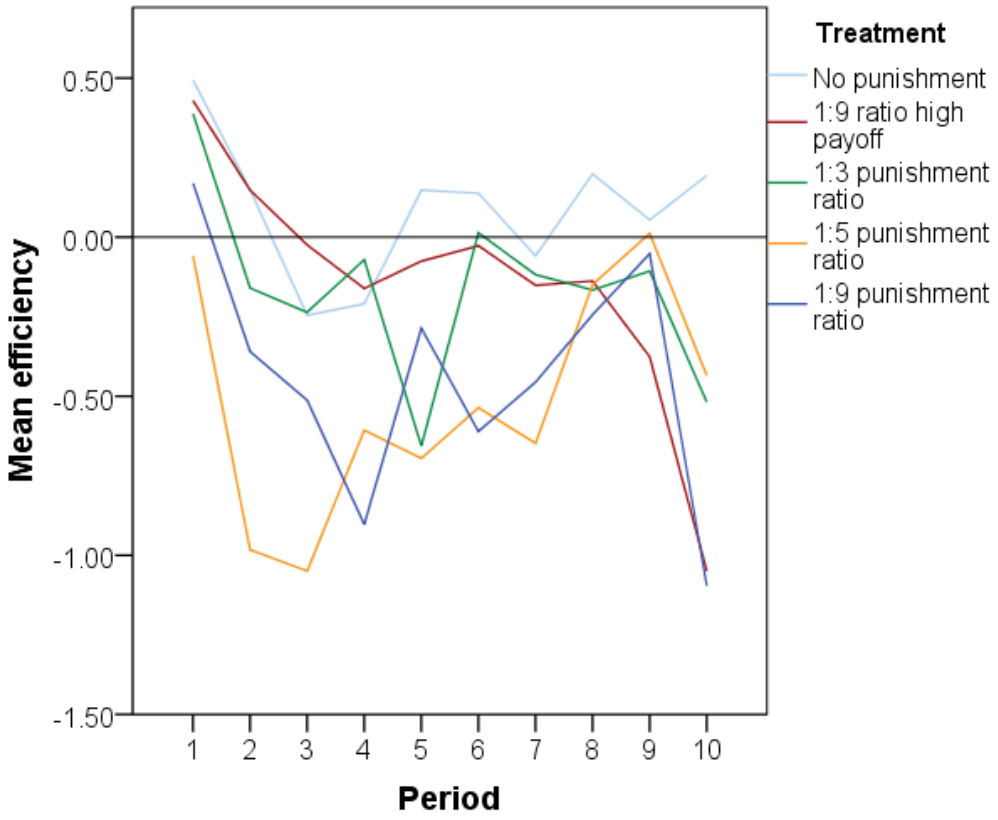

Table 17: Mean efficiency

Table x shows mean efficiency levels with standard errors in parentheses. P-values for Wilcoxon signed-rank tests of difference between the first and last five periods are in the right column.

	Periods 1-10	Periods 1-5	Periods 6-10	Wilcoxon signed-rank p-value
No punishment	0.09 (1.05)	0.07 (1.27)	0.11 (0.78)	0.704
1:3 punishment	-0.16 (1.13)	-0.15 (1.19)	-0.18 (1.08)	0.679
1:5 punishment	-0.52 (1.20)	-0.68 (1.25)	-0.35 (1.13)	0.022
1:9 punishment	-0.44 (1.42)	-0.38 (1.46)	-0.49 (1.38)	0.209
1:9 high payoff	-0.14 (1.11)	0.06 (1.11)	-0.35 (1.08)	0.002

Table 18: Differences in efficiency

P-values for Mann-Whitney tests (2-tailed) of difference in efficiency levels between punishment treatments.

Periods 1-10

	1:3 punishment ratio	1:5 punishment ratio	1:9 punishment ratio	1:9 high payoff
No punishment	0.001	0.000	0.000	0.005
1:3 punishment	.	0.000	0.048	0.608
1:5 punishment	.	.	0.032	0.000
1:9 punishment	.	.	.	0.012

Periods 1-5

	1:3 punishment ratio	1:5 punishment ratio	1:9 punishment ratio	1:9 high payoff
No punishment	0.065	0.000	0.002	0.849
1:3 punishment	.	0.000	0.140	0.063
1:5 punishment	.	.	0.001	0.000
1:9 punishment	.	.	.	0.001

Periods 6-10

	1:3 punishment ratio	1:5 punishment ratio	1:9 punishment ratio	1:9 high payoff
No punishment	0.003	0.000	0.000	0.000
1:3 punishment	.	0.317	0.227	0.243
1:5 punishment	.	.	0.817	0.811
1:9 punishment	.	.	.	0.786

5. Discussion

5.1 Levels of cooperation and extraction

At a high level, this research finds that behaviour in CPR games responds to incentives and punishment in a very similar way to the linear VCM game. In particular, incentives to defect or cooperate (defined by the payoff functions) have a strong effect on the observed level of cooperation, as does the level of punishment that occurs. Further, punishment behaviour is responsive to price and income. Our hypothesis was that stronger punishment can improve cooperation in economic environments with higher incentives to defect, such as is observed within CPR games. The significant reductions in extraction resulting from increased punishment effectiveness ratios in our design show this hypothesis to be correct. There is strong statistical evidence to reject the null hypothesis that the strength of punishment has no effect on extraction levels.

This finding is a positive outcome in terms of implications for the management of nitrate as a common pool resource. Even though the incentives to defect are strong and coupled to profit, the responsiveness to changes in incentives and punishment effectiveness we observe in this study implies that getting the institutional settings right in relation to economic incentives can significantly improve the chances of success. One further implication is that any group property approach to managing nitrate losses will need to explicitly consider volatility, both in terms of how production is affected by the weather and how commodity prices affect the relationship between profit and production. Volatility means that the incentives to cooperate or defect will change from year-to-year. Given that behaviour is responsive to changes in the incentive environment, institutional settings will need to explicitly consider this.

Our results are consistent with those of Ostrom et al. (1992), in that a relationship was evident between the cost of punishment, demand for punishment and levels of cooperation. Stronger punishment is clearly able to induce higher levels of cooperation. Extraction levels decrease monotonically with respect to more

effective, lower-cost punishment. This finding suggests that the lesser effectiveness of peer punishment that has been observed in previous experiments in CPR games is not due only to confusion. The incentive environment and the cost of punishment also play a part.

We appear to have a significantly higher propensity for defection in our population than that of Cason and Gangadharan. This is an interesting result given the use of identical incentive environments, punishment technologies and procedures. While still finding a punishment ratio of 1:3 to result in extraction levels significantly below those without punishment, even at a punishment ratio of 1:9, our extraction levels were still slightly higher than those of Cason and Gangadharan with a punishment ratio of 1:3.

It is also notable that we appeared to have correspondingly greater levels of decline in cooperation over time. Cason and Gangadharan (2015) reported a decrease in cooperation over time in the case of their treatment with no punishment, and no breakdown over time in their 1:3 punishment treatment. Our no punishment treatment started with high contributions, but stayed at about the same level (close to the Nash equilibrium) over time, with a Spearman p-value of 0.215 (as shown in tables 10 and 12). Our 1:3 punishment ratio treatment showed a clear decline, with a Spearman p-value of 0.008. Increasing the ratio to 1:5 and then 1:9 led to a flatter response evident in this curve. Given the strong evidence of profitable deviations in our experiments with weak punishment (as shown in figures 12 and 13), it is perhaps not surprising that we see this pattern.

Introducing peer punishment to the CPR environment has some capacity to improve cooperation for some groups, at any level of punishment effectiveness ratio. Even weak punishment can induce a response from some participants by acting as a social signal. The data show some evidence of "crowding out"; in that increasing the effectiveness of punishment to 1:5 induced a lesser responsiveness to punishment, with the increase to 1:9 counteracting this by removing the potential for defection to pay off. With weak punishment (and indeed, even in the absence of punishment), social preferences alone, coupled with the disapproval signalled by the use of weak punishment, is sufficient to support cooperation for

many groups. However, this cooperation rapidly breaks down in the presence of defectors or participants with more self-interested motivations.

Increasing the effectiveness of peer punishment and decreasing its cost had significant positive effects on extraction levels, both in terms of bringing them closer to the social optimum and in terms of making extraction levels more stable over time. The implication for our theoretical common-pool resource is positive, in that increased punishment results in environmental improvements. However, extraction levels were still significantly above the social optimum with a mean of 6.86, while the social optimum was 5. This occurred despite the fact that, on average, increased defection from the group average resulted in lower income. There are two potential explanations worthy of further consideration here: hedging and confusion.

Hedging

Even among groups with relatively high levels of cooperation (in terms of the level of agreement between the group members, as evidenced by their contributions), contributions were often still slightly above the social optimum. This same effect was observed by Cason and Gangadharan (2015)

> *In particular, for the CPR treatment with punishment, only 10 percent of the individual extraction levels were at the social optimum of 5. The most frequent extraction was 6, chosen 35 percent of the time, and the modal extraction was 6 in every period. The data do suggest some failed attempts to coordinate on the social optimum of 5, since the fraction of choices on 5 falls in half from periods 1-4 to periods 5-10. Although cooperation on an extraction level of 6 rather than 5 is suboptimal, it permits subjects to earn nearly the same level of profit (50.4) as the 52 earned at the social optimum. Importantly, this suboptimal extraction level of 6 substantially reduces the incentives to deviate, typically by 10 to 12 experimental dollars* (Cason and Gangadharan, 2015).

Parameterised as it is, the marginal incentives for defection in this CPR game are very high. Overall outcomes are extremely sensitive to even small deviations from the optimum as a result. This sensitivity makes coordination particularly difficult, as there are large consequences for small failures. It is possible that by

contributing 6 tokens, groups that are generally cooperative are hedging against damage from defection. However, this is not the only possible explanation.

Confusion

It is important to note that the CPR game is, fundamentally, very difficult. Accordingly, it is quite possible that the choice of 6 out of 12 tokens as an extraction level could reflect a small level of confusion, for participants that have failed to optimise returns across both the private and public accounts. In calculating marginal returns from the public account against marginal returns from the private account, the optimal number of tokens in the group account is 20, or 5 tokens for each group member. However, the maximum payoff from the group account occurred at 22 tokens. Participants can easily make the mistake of the optimum being 5-6 tokens, if they are not paying careful attention to the marginal returns from the private account as well. This feature is non-trivial. Indeed, while we do not have formal survey evidence to support this, many economics students reported the optimum being "5 or 6" tokens, when asked if they had a strategy on exit. A valuable extension to this study would be to present a sizeable sample from the same population with the payoff environment and query what they believed to be the socially optimal decision. If it is indeed the case that many participants are mistaken about the optimum, then the result of the mean extraction level of 6.86 in the treatment with a punishment ratio of 1:9 is, in fact, quite good.

There is one further problem of confusion that could cause the observed frequency of extraction levels of 6 tokens. Simply put, participants may be so confused as to be unable to apply any clear strategy. In this case, participants may be simply splitting their tokens evenly between the private and public accounts. Again, this possible explanation is well worth investigating with survey evidence. In either the case of hedging or confusion, these possibilities emphasise the importance of resource users being able to communicate in order to clarify the best actions for both profit and environmental management.

5.2 Incentives

Weak punishment is not very effective in improving cooperation for common pool resource management in the presence of strongly self-interested participants. If defectors find that their earnings exceed any damage they receive as the result of punishment, they will quickly learn to extract ever-higher levels. Self-interested individuals are able to determine that the costs of defection (even when receiving heavy punishment) are easily outweighed by the benefits from increased earnings. This was particularly evident for participants who made large earnings early in the game. By amassing a stockpile of earnings early on, these participants were free to pursue aggressive and risky strategies later in the game, safe in the knowledge that they already had significant earnings. This is similar to effects observed in industrial organisation games. For example, if an asymmetry in the market develops early on, richer or larger firms are able to operate more competitively (even operating at a loss) in order to defeat their competition using predatory pricing (Harrison, 1988). This is an important implication for the potential management of nitrate losses as a common pool resource. While communication and social capital might be effective with groups that have come together with a sense of common purpose and an environmental goal, this could be unstable over time. Changes in group composition may mean that more self-interested individuals begin to exploit institutional arrangements if they are not designed with the incentive environment in mind. If payoffs from defection cannot be counteracted by sanctions in this context, cooperation may begin to break down.

Rivalry in the CPR environment means that behaviour is often even more aggressive than predicted by game theory (Van Soest & Vyrastekova, 2006). In some cases, defectors also punished co-operators heavily (for example, extracting the maximum 12 units from the CPR as well as assigning five punishment points to every other player). Not only did these defectors have ample income with which to purchase punishment, but the effect of punishment on lower earners was proportionally much worse. As the co-operators had much lower earnings, the effect of punishment was severe for them, leaving them with very little cumulative earnings. Being left with very low earnings may discourage these

individuals from spending money to punish defectors in subsequent rounds. The punishment from defectors sometimes had the effect of driving co-operators' extraction levels down even lower.

A brief discussion of the different payoffs associated with different strategies is useful here in developing an understanding of the incentives associated with these different actions, and the dramatic effects on earnings of the different extraction and punishment strategies. As covered in section 3.6 regarding game theoretic outcomes, the socially optimal contribution in our low payoff incentive environment is 5 tokens from each participant, giving them a payoff of 52 experimental dollars each. If one individual defects from this level and contributes 12, they receive a payoff of 74.4, reducing the payoff for the co-operators to 38. Given that the Nash strategy for this game is to extract 8 from the CPR, this is an aggressive strategy that reflects an expectation that the others will cooperate, enabling asymmetric gains to be captured by the single defector.

Earnings interact with punishment effectiveness. Assume we use the punishment ratio of 1:3 that is commonly used in the VCM literature. If the defector in this case assigns a full five deduction points to all other participants, their individual payoffs are reduced by 15, leaving them with 23 experimental dollars from that round. However, even if all three other participants punish the one defector the full amount possible, the defector will only receive a deduction of 45 experimental dollars, leaving the defector with a payoff of 29.4, which is still 6.4 experimental dollars better off than if they had cooperated. Even this level of difference requires all three defectors acting in concert to punish the defector, which is extremely unlikely without communication, given that punishment itself is a second-order public good that is subject to free riding in the same manner as extraction levels (Chaudhuri, 2011). Clearly, the punishment effectiveness ratio of 1:3 is likely to be woefully inadequate to discourage defection in this incentive environment.

Increasing the punishment ratio to 1:5 for this scenario, the payoffs of 74.4 and 38 for defector and co-operator (respectively), would be reduced to -0.6 and 13 if the same punishment strategies were pursued as in the previous example. Increasing

the ratio to 1:9 would leave payoffs at -60.6 and -7. Of course, this does not occur in practice due to free-riding on punishment. Importantly though, at this level of punishment, a single participant is able to destroy all the defector's gains, reducing their earnings to 29.4 experimental dollars. Even with free riding from two other participants, the defector can be discouraged. Given that punishment was usually targeted towards defectors, this explains to a certain extent why defectors did not on average earn better than co-operators in this treatment.

In practice, the strategies of co-operators in response to a defector were mixed. Insufficient punishment to reduce the earnings of the defector was common. Interestingly, many participants responded to high extraction from other group members by reducing their own extraction. If their assumption was that the other two participants would continue to cooperate, the optimal extraction level remains 5 for these three, while the Nash strategy increases slightly from 8 to 9 tokens. Contributions below this level either indicate confusion, or a desire to produce optimal returns from the CPR regardless of the distribution.

As an extreme example, if a defector extracts 12 and all other participants extract 0, the earnings of the defector become 146.4 compared to the earnings of 12 from other participants. While this might seem unlikely, the highest earnings before punishment of any participant were 103.2. This is well in excess of what might be earned without some participants reducing their extraction levels to compensate for defection impacts on the total CPR earnings. Clearly even the punishment ratio of 1:9 is not sufficient to remove the gains from defection in all cases, though it was sufficient to do so on average within the results presented in Section 5 (for example, as evidenced in figure 12).

5.3 Punishment

The cost of punishment

As for Ostrom et al. (1992), Carpenter (2002a), Andreoni et al. (2003), Putterman & Anderson (2003) and Nikiforakis & Normann (2008), we found the demand for punishment to be inversely related to its cost and targeted mostly towards defectors. By decreasing the cost of punishment, we were able to generate

increased demand for punishment. This increased punishment became sufficient to remove the ability of defectors to earn more than co-operators, on average. Punishment (particularly at the 1:9 level in our low payoff environment) was affordable enough for those with low earnings and effective enough to remove the advantage accruing to defectors, on average.

The decrease in cost from punishment effectiveness ratios of 1:3 to 1:5 generated a large demand response, but the increase in punishment used from 1:5 to 1:9 was much more modest, as shown in the demand curve in figure 9. Further increasing the punishment effectiveness ratio is likely to increase demand for punishment further, but only by relatively small amounts if the general shape of the demand curve is similar to what we have observed. This is consistent with the diminishing marginal returns to punishment evident in the experimental results reported in Section 4. In general, the incentives (particularly marginal incentives) to defect are much higher in CPR games. Accordingly, punishment technologies designed to counter the incentives present in the linear VCM game are wholly inappropriate in this environment.

Income effects

Changes to the payoff environment interacted dynamically with the cost of punishment. Increased incomes in our high payoff environment drove much higher demand for punishment. However, the income elasticity of demand for punishment was low at 0.66. This meant that the increase in punishment used by participants in the game was insufficient to counteract the effect of much higher earnings for defectors, even though the percentage gain for cooperation was held constant.

The increased earnings for defectors in our high payoff environment drove a pronounced decline in cooperation over time. While the earnings of co-operators also increased, they did not choose to spend proportionally more on punishment. Because of this dynamic interaction between incomes, the quantity of punishment demanded and the level of defection that occurred, the effect of income on punishment behaviour or overall cooperation is likely to be difficult to predict

without rigorous evaluation within a specific incentive environment involving a specific population. In the context of agricultural nitrate losses, information about the abatement curves of typical farms in the area would be useful to inform discussions of what sanctions might be appropriate among users.

5.4 Relationship between incentives, punishment ratio and cooperation

The appropriate punishment ratio for a given incentive environment is a product of the interaction between the use of punishment and its impact on defectors. The non-linear incentives of the CPR environment in particular make it difficult to establish a clear relationship between economic incentives and the required level of punishment ratio except through experimentation. We began with two potential explanations for an effective punishment ratio:

1. Punishment levels respond to damage levels: where cooperators' willingness to use punishment is related to the amount of damage that they receive, due to the gains of defectors. This behaviour would indicate a greater willingness to use punishment, in order to counteract greater damage caused by defection.
2. Prospering cheaters: where demand for punishment is not directly related to damages or inequality aversion, and that the likelihood of defection leading to high returns being the primary driver of punishment effectiveness.

In order to discriminate between these two potential drivers, our design used punishment ratios of 1:3, 1:5 and 1:9, in order to provide participants with punishment technologies of similar effectiveness to those that are successful in the linear VCM, with each ratio assuming a different relationship between the incentive environment and the effectiveness of punishment.

While there are strong incentives to defect in our CPR games, pursuit of these gains by defectors generates significant damage for co-operators, in that it greatly reduces their potential earnings. We hypothesised that this may induce a response to damages effect, where co-operators choose to expend more money on

punishment in order to avoid damages from over-extraction by defectors. In our low payoff CPR games, we found that the ratio of 1:5, based on the maximum gains and damages achieved in the CPR game (relative to a 1:3 ratio in the VCM game) had some effectiveness in reducing over-extraction. However, this was not significantly different to our treatment with a 1:3 punishment ratio.

Within the normal payoff CPR environment, the punishment ratio of 1:9 has two possible relationships to a ratio of 1:3 in the VCM environment. First, 1:9 relates to the maximum marginal gains to marginal damages ratio (the point where maximum payoffs from defection and maximum damages to co-operators coincide). Second, it relates to the highest possible payoff through defection (the maximum gains for defectors that would need to be destroyed in order to discourage defection). This ratio was much more successful, generating significantly lower extraction levels and an extremely flat level of extraction over time. Notably, a punishment ratio of 1:9 was able to make the incomes of defectors lower than those of co-operators, on average.

Given the two possible explanations for the success of this punishment ratio, one further treatment was needed to differentiate between the two. In our high payoff environment, a punishment ratio of 1:9 is in line with the marginal gains to damages explanation. An explanation related to the highest payoff from defection would suggest that a higher punishment ratio of 1:14 would be required. We found that a punishment ratio of 1:9 was ineffective in the high payoff environment. A marked decline in cooperation over time was evident, which was statistically significant.

Given this combination of findings, the response to damages effect does not appear to be present. This is consistent with the findings of Falk, Fehr and Fischbacher (2005) in linear VCM games. Changing the experimental framing to provide feedback to participants that includes the stage one earnings of every player, as well as their contributions, may induce more punishment from co-operators. That is, it is possible that we cannot expect to see decisions based on inequality-aversion when we do not communicate to participants the full state of inequality that exists. In the field, the level of wealth accrued from defection may

be the more obvious indicator of defection than the level of extraction, this is worth consideration in future experimental evaluations, especially those to be carried out in the field.

The more important driving factor behind what punishment ratio is effective would seem to be whether cheaters prosper. In order to sustain cooperation over the long term, the level of punishment that occurs must reach a high enough level that defectors do not earn more than co-operators, as seen in our 1:9 punishment ratio treatment for the standard payoff incentive environment. In institutional design for resource management in the field, the cost and effort required to sanction defectors to a point where they do not prosper from defection should therefore form a useful basis for discussion about the design of robust institutions. However, determining this relationship is not likely to be easy or straightforward.

Group members do not select punishment levels by calculating their own damages relative to the earnings of others. Rather, they select punishment based on what they feel is worthwhile (in terms of its effectiveness relative to cost) and affordable (in terms of their income). Response to punishment, however, tends to be more directly connected to incentives. It is not until the impacts of punishment exceed the potential benefits from defection that we see the decline in cooperation over time fully abate. Accordingly, selection of an appropriate fine level for environments featuring peer punishment is not straightforward. Participants' demand for punishment was inelastic. This indicates that punishment effectiveness ratios must increase more than the incentives to defect in a given environment. This inelastic relationship means that managing changes in incentives caused by exogenous factors (such as weather and commodity price volatility) becomes more challenging. If it was a perfect 1:1 relationship, we could expect resource users to manage defection with static tools, as an increase in incentives to defect would have a corresponding increase in willingness to use punishment. The absence of this effect means it is important that institutions explicitly consider volatility.

Figure 15 shows the interaction between different experiment parameters and behaviour that was observed in this study. Punishment cost had a large, nonlinear

impact on the overall amount of punishment. The payoff environment affected the amount of punishment, with an income elasticity of 0.66. The payoff environment changed the incentives to defect, with a corresponding positive impact on the amount of extraction observed. Finally, the amount of punishment that occurred (as driven by both cost and payoffs) had an impact on the amount of extraction that occurred.

It is possible that a theoretical relationship could be derived for these interactions, but the amount of data required is non-trivial. Though this study did not produce sufficient data to develop a quantitative model for the relationship between punishment cost, payoffs and extraction, the data clearly support some form of dynamic interaction that is defined by the relationships of influence represented by figure 15.

Figure 15: Influence diagram for impacts of punishment effectiveness

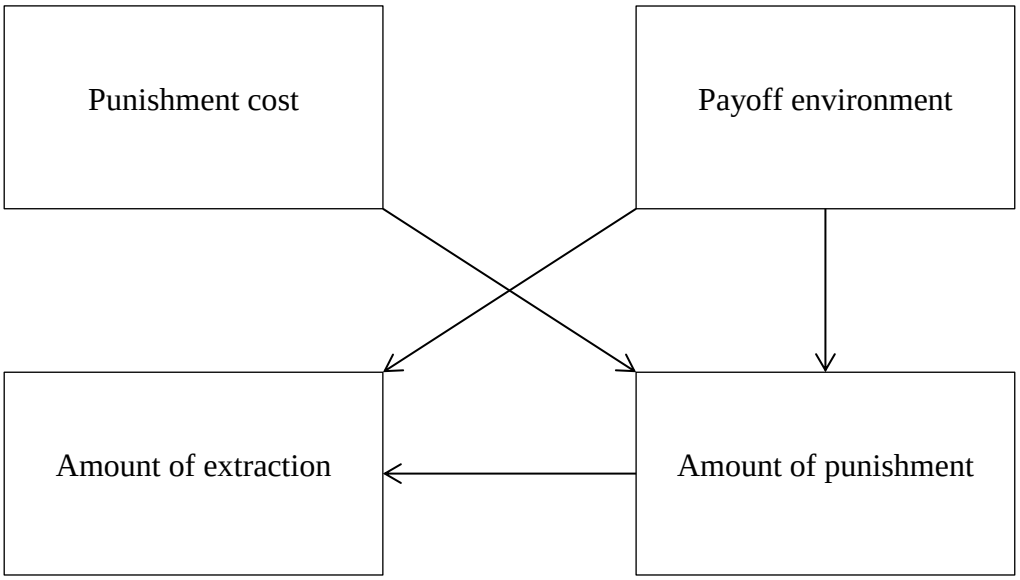

5.5 Efficiency

Of course, reducing resource extraction is not the only concern in effective management of common pool resources. The overall level of welfare accrued to a group of resource users from their resource use and institutional constraints is also

of interest. Enhancing efficiency reflects balancing the twin goals of achieving environmental outcomes while still enabling profitable farming. Our findings in relation to efficiency are clearly in line with the findings of Ostrom et al (1992), in that "swords without a covenant may be worse than the state of nature". While we were able to significantly improve the extraction levels for our common pool resource, the overall welfare effect of including punishment was negative. As for Janssen, Holahan, Lee, & Ostrom (2010), all of our treatments had initial declines in efficiency, followed by gradual improvement. Spiteful punishment in the final period had a large effect on final period efficiency. On the whole, efficiency was poor, with all treatments that included punishment having worse total welfare outcomes than the treatment without punishment.

While the findings for efficiency in this study were poor, it is possible that more periods would enable this to improve (Gächter, Renner, & Sefton, 2008). The increased complexity of the CPR game compared to the linear game clearly makes determining an optimal contribution level more difficult. This increased complexity is also likely to inhibit learning, slowing participants' response to feedback (Bereby-Meyer & Roth, 2006). First, interpreting the contributions of others would be more difficult. Second, understanding the reasoning behind the punishment behaviour of others would be more difficult. Participants that are repeatedly punished for their behaviour, but with little understanding of the reasons why this is done, are likely to experience a form of learned helplessness, wherein they default to repeating a passive strategy (such as undertaking zero extraction).

5.6 Implications for collective management of water quality

This research is a positive story for self-governance of the commons. Our experimental results show clearly that members of a group faced with a common pool resource dilemma respond to incentives in the same way that is expected for other types of social dilemma. In the context of collective management of nitrate losses, this means institutional design should consider the incentives to defect, and the costs of effective punishment to counteract this. With design taking careful

account of the abatement costs and profit functions of typical farms in a particular area, the likelihood of successful management will be enhanced.

Communication and social capital will still have an important role to play, but this should perhaps be considered differently than purely in terms of effects on cooperation. Given that low-cost, strong punishment can be effective in preventing declines in cooperation over time, the role of communication becomes primarily one of supporting efficiency. The large degree of complexity involved in managing common pool resources means that communicating about the right actions to take (and when) is important to establishing agreed norms for resource use. In the field, this is likely to be further complicated by the introduction of heterogeneity, as norms for resource use may not necessarily be exactly the same for different types of operation. Volatility caused by weather and commodity prices will also have implications for the role of communication. As the incentives change from one year to the next, maintaining a shared understanding of norms through on-going engagement is likely to be important.

If preventing a decline in cooperation over time were wholly dependent on good communication, then we might expect community governance to be fragile. Inevitably, community groups will experience breakdowns in communication, changes to the group composition and other disruptions. If communication were critical to sustaining resource use levels, then these disruptions would become perilous for maintaining the resource. However, if communication is primarily supporting efficiency, then temporary disruptions to this part of the institution are not likely to be fatal. Effective sanctions and rules could potentially operate in the absence of good communication for a time, with only a temporary implication for efficiency.

This being said, given the great expense associated with the use of sanctions without communication, it is clear that this would not be a desirable arrangement for governing a common pool resource. In the interest of welfare maximisation, the use of communication and threats are essential to avoid the significant costs of using costly sanctions. The combination of communication and sanctions is likely to better support meeting environmental outcomes and maintaining profitable farms. The task of optimising resource use in our laboratory environment is

difficult, generating significant amounts of confusion. It would be hard to overstate the greater level of complexity that is associated with managing the resource effectively in the field, compared to our stylised representation of the problem in the laboratory. Communicating and working together to develop shared understanding of the twin problems of optimising farm profit while managing a cumulative impact on a common resource is essential.

6. Summary and Conclusion

6.1 The relationship between incentives, punishment and cooperation in common pool resource games

In our experiments, a clear relationship was shown between the strength of punishment and the level of resource extraction that occurred. Participants responded clearly to incentives, even in the absence of communication. While the highly-complex environment of the common-pool resource experiment makes coordination difficult without communication, this research shows that it can clearly be improved.

Punishment effectiveness (and its cost) has a strong effect on the amount of punishment that is demanded by individuals within the CPR game. In addition, the income of individuals affects the quantity of punishment that they purchase. If this level of punishment is able to remove the ability of participants to earn more through defection than they would through cooperation, levels of resource use can be sustained over time that are well-below the Nash equilibrium.

The incentive structure of the economic environment also has a profound effect on whether cooperation can be sustained. While raising incomes can increase the quantity of punishment people are willing to purchase, it also increases the income from defection. The propensity to increase punishment does not increase at the same rate as the payoffs from defection. This means that higher stakes require stronger punishment ratios.

If the dynamic interaction between the incentive environment and the cost of punishment can be managed to provide a sufficient quantity of punishment to reduce the earnings of defectors, cooperation can be sustained over time. However, this may still be at sub-optimal levels due to hedging against defection or confusion about the optimal level of extraction.

6.2 Limitations and future extensions

The primary limitation of this study has been the failure to provide an explanation for the continued gap between the optimal level of extraction and the observed level of extraction, even in highly cooperative groups with stable levels of extraction. The possible explanations of hedging or confusion both remain extant. Survey work would be valuable in disentangling these two factors. Directly informing participants of the social optimum and Nash equilibrium levels of extraction would potentially provide a strong priming effect, but it is questionable whether it would have a much different effect relative to the linear VCM setting, given that the optimal contribution levels are much more obvious to participants in that context. In any case, in the field it is likely that decisions about allocation and extraction levels would be made ex-ante, making navigation of the complexity of optimisation much more straightforward for resource users.

Communication is well known to improve levels of cooperation. Given that an environment with no communication is unlikely to be found in the field, this makes the experiments slightly artificial. It would be an interesting extension to repeat the same experiments with the inclusion of communication, to determine whether altering the punishment ratio still has any significant effect on the overall levels of cooperation observed. Perhaps limited or imperfect forms of communication could bring the scenario closer to that observed in the field.

It is a notable feature of the punishment technology that we have used that as the punishment ratio increases, the tool becomes increasingly coarse, as punishment increases in ever-larger increments. It may be worthwhile to consider punishment technologies that enable smaller increments at the same cost.

Finally, the realism of the experimental environment could be greatly improved by the introduction of uncertainty or variability. As detailed in the background section, non-point source pollution is characterised by both uncertainty and variability. Variations in payoff functions from period-to-period would provide interesting insights as to how cooperation is affected by under- or over-sensitivity to rare events. A more ambitious design (likely requiring many more periods to

enable sufficient time for learning) could also include incomplete information about payoffs in order to investigate the experience-description gap. This provides more realism, but also complicates the decision task in what is already a complex experimental environment.

6.3 Concluding remarks

In the presence of some form of external threat or regulation associated with the degradation of water quality, the diffuse loss of nitrate to waterways is transformed into a common pool resource problem. This social dilemma is characterised by strong incentives to defect and challenging levels of complexity. However, these challenges are not insurmountable.

Due to the complexity of common pool resource management, communication and agreement about norms for resource use levels is important to support groups in optimising environmental outcomes and profitability. Communication is a necessary condition for maximising welfare when navigating complexity. However, communication is not necessary to improve environmental outcomes or sustain a level of resource use over time.

Sanctions are a critical part of the sustainable management of common-pool resources. While some self-selected groups of highly cooperative individuals may be able to sustain cooperation in the absence of sanctions, these arrangements are unlikely to be stable over time. The presence of defectors means that some form of recourse to punishment is necessary to support the continued sustainable management of a resource. This has been well-established through data from both laboratory experiments and the field (Ostrom, 1999). The strong incentives for defection in common pool resource management, including collective management of water quality, can be overcome by sufficiently strong sanctions.

Our experimental findings suggest that sanctions administered by peers can be effective in improving environmental outcomes and stabilising the level of resource use over time, preventing the decline in cooperation that is in evidence

where individuals are able to profit from defection. The cost and effectiveness of sanctions has a strong interaction with the nature of the incentive environment. In designing institutions for the collective management of nitrate loads to water bodies, regulators and groups of resource users must consider the incentives for defection that are driven by the degree of coupling of nitrate losses and profitability in farm systems. In order to be successful, group members must have access to low-cost sanctions that are designed with the incentive environment in mind. If institutions can be designed in such a way as to ensure that cheaters do not prosper from defection, the prospects for collective management of water quality are positive.

7. References

Adger, W. N., & Luttrell, C. (2000). Property rights and the utilisation of wetlands. *Ecological Economics, 35*(1), 75-89.

Ambrus, Attila, & Greiner, Ben. (2012). Imperfect Public Monitoring with Costly Punishment: An Experimental Study. *American Economic Review, 102*(7), 3317-3332.

Aminova, M., & Abdullayev, I. (2009). Water management in a state-centered environment: water governance analysis of Uzbekistan. *Sustainability, 1*(4).

Anderson, Christopher M., & Putterman, Louis. (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), 1-24.

Andreoni, James. (1988). Why free ride?: Strategies and learning in public goods experiments. *Journal of Public Economics, 37*(3), 291-304.

Andreoni, James, & Gee, Laura K. (2012). Gun for hire: delegated enforcement and peer punishment in public goods provision. *Journal of Public Economics, 96*(11), 1036-1046.

Apesteguia, Jose, & Maier-Rigaud, Frank P. (2006). The role of rivalry public goods versus common-pool resources. *Journal of Conflict Resolution, 50*(5), 646-663.

Athanassoglou, Stergios. (2010). Dynamic nonpoint-source pollution control policy: ambient transfers and uncertainty. *Journal of Economic Dynamics and Control, 34*(12), 2494-2509.

Ballantine, D. J., & Davies-Colley, R. J. (2013). Water quality trends in New Zealand rivers: 1989--2009. *Environmental Monitoring and Assessment, 186*(3), 1939–1950.

Balliet, Daniel. (2010). Communication and Cooperation in Social Dilemmas: A Meta-Analytic Review. *Journal of Conflict Resolution, 54*(1), 39-57.

Barclay, Pat. (2006). Reputational benefits for altruistic punishment. *Evolution and Human Behavior, 27*(5), 325-344.

Bay of Plenty Regional Council. (2012). Lake Rerewhakaaitu. Retrieved from http://www.boprc.govt.nz/media/33637/RotoruaLakes-100827-LakeRerewhakaaituReport.pdf

Ben-Ner, Avner, & Putterman, Louis. (2009). Trust, communication and contracts: An experiment. *Journal of Economic Behavior & Organization, 70*(1), 106-121.

Bereby-Meyer, Yoella, & Roth, Alvin E. (2006). The Speed of Learning in Noisy Games: Partial Reinforcement and the Sustainability of Cooperation. *The American Economic Review, 96*(4), 1029-1042.

Blau, Peter Michael. (1964). *Exchange and power in social life*: Transaction Publishers.

Bochet, Olivier, Page, Talbot, & Putterman, Louis. (2006). Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior & Organization, 60*(1), 11-26.

Bochet, Olivier, & Putterman, Louis. (2009). Not just babble: Opening the black box of communication in a voluntary contribution experiment. *European Economic Review, 53*(3), 309-326.

Bohm, Peter. (1972). Estimating demand for public goods: An experiment. *European Economic Review, 3*(2), 111-130.

Bohnet, Iris, Frey, Bruno S., & Huck, Steffen. (2001). More Order with Less Law: On Contract Enforcement, Trust, and Crowding. *The American Political Science Review, 95*(1), 131-144.

Bosch, Darrell J., & Pease, James W. (2000). Economic Risk and Water Quality Protection in Agriculture. *Review of Agricultural Economics, 22*(2), 438-463.

Bowles, Samuel. (2008). Policies Designed for Self-Interested Citizens May Undermine "The Moral Sentiments": Evidence from Economic Experiments. *Science, 320*(5883), 1605-1609.

Bowles, Samuel, & Gintis, Herbert. (2002a). Homo reciprocans. *Nature, 415*(6868), 125-128.

Bowles, Samuel, & Gintis, Herbert. (2002b). Social Capital and Community Governance. *The Economic Journal, 112*(483), F419-F436.

Bowles, Samuel, & Polanía-Reyes, Sandra. (2012). Economic Incentives and Social Preferences: Substitutes or Complements? *Journal of Economic Literature, 50*(2), 368-425.

Burger, J., & Gochfeld, M. (1998). The tragedy of the commons 30 years later. *Environment: Science and Policy for Sustainable Development, 40*(10), 4-13.

Burton, Rob J. F., & Paragahawewa, Upananda Herath. (2011). Creating culturally sustainable agri-environmental schemes. *Journal of Rural Studies, 27*(1), 95-104.

Cabe, Richard, & Herriges, Joseph A. (1992). The regulation of non-point-source pollution under imperfect and asymmetric information. *Journal of Environmental Economics and Management, 22*(2), 134-146.

Carpenter, Jeffrey P. (2007a). The demand for punishment. *Journal of Economic Behavior & Organization, 62*(4), 522-542.

Carpenter, Jeffrey P. (2007b). Punishing free-riders: How group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior, 60*(1), 31-51.

Casari, Marco. (2005). On the Design of Peer Punishment Experiments. *Experimental Economics, 8*(2), 107-115.

Casari, Marco, & Luini, Luigi. (2012). Peer punishment in teams: expressive or instrumental choice? *Experimental Economics, 15*(2), 241-259.

Casari, Marco, & Plott, Charles R. (2003). Decentralized management of common property resources: experiments with a centuries-old institution. *Journal of Economic Behavior & Organization, 51*(2), 217-247.

Cason, Timothy N., & Gangadharan, Lata. (2015). Promoting cooperation in nonlinear social dilemmas through peer punishment. *Experimental Economics, 18*(1), 66-88.

Cason, Timothy N., & Gangadharan, Lata. (2013). Empowering neighbors versus imposing regulations: An experimental analysis of pollution reduction schemes. *Journal of Environmental Economics and Management, 65*(3), 469-484.

Cason, Timothy N., & Gangadharan, L. (2016). Swords without covenants do not lead to self-governance. Journal of Theoretical Politics, 28(1), 44-73.

Chaudhuri, Ananish. (2011). Sustaining cooperation in laboratory public goods experiments: a selective survey of the literature. *Experimental Economics, 14*(1), 47-83.

Choi, Jung-Kyoo, & Ahn, T. K. (2013). Strategic reward and altruistic punishment support cooperation in a public goods game experiment. *Journal of Economic Psychology, 35*, 17-30.

Clapcott, Joanne, Young, Roger, Goodwin, Eric, Leathwick, John, & Kelly, David. (2011). Relationships between multiple land-use pressures and individual and combined indicators of stream ecological integrity *DOC Research and Development Series*. Wellington: Department of Conservation.

Cochard, François, Willinger, Marc, & Xepapadeas, Anastasios. (2005). Efficiency of Nonpoint Source Pollution Instruments: An Experimental Study. *Environmental and Resource Economics, 30*(4), 393-422.

Coglianese, Cary, & Lazer, David. (2003). Management based regulation: Prescribing private management to achieve public goals. *Law & Society Review, 37*(4), 691-730.

Coleman, James S. (1990). Foundations of social capital theory. *Cambridge, Mass: Belknap*.

Coleman, James S. (1993). The Rational Reconstruction of Society: 1992 Presidential Address. *American Sociological Review, 58*(1), 1-15.

Cox, Michael, Arnold, Gwen, & Villamayor-Tomás, S. (2010). A review of design principles for community-based natural resource management. *Ecology and Society, 15*(4), 38.

Cox, Michael, & Ross, Justin M. (2011). Robustness and vulnerability of community irrigation systems: The case of the Taos valley acequias. *Journal of Environmental Economics and Management, 61*(3), 254-266.

Cudney-Bueno, R., & Basurto, X. (2009). Lack of cross-scale linkages reduces robustness of community-based fisheries management. *PLoS ONE*(July), e6253.

Dawes, Robyn M., McTavish, Jeanne, & Shaklee, Harriet. (1977). Behavior, communication, and assumptions about other people's behavior in a commons dilemma situation. *Journal of Personality and social Psychology, 35*(1), 1-11.

Dodd, Mike, Wilcock, Bob, & Parminter, Terry. (2009). Review of recent rural catchment-based research in New Zealand: AgResearch.

Doole, G. J. (2010). Indirect instruments for nonpoint pollution control with multiple, dissimilar agents. *Journal of Agricultural Economics, 61*, 680-696.

Doole, G. J. (2012). Cost-effective policies for improving water quality by reducing nitrate emissions from diverse dairy farms: An abatement–cost perspective. *Agricultural Water Management, 104*, 10-20.

Doole, G. J., & Kingwell, Ross. (2015). Efficient economic and environmental management of pastoral systems: Theory and application. *Agricultural Systems, 133*, 73-84.

Doole, Graeme J., & Pannell, David J. (2012). Empirical evaluation of nonpoint pollution policies under agent heterogeneity: regulating intensive dairy production in the Waikato region of New Zealand*. *Australian Journal of Agricultural and Resource Economics, 56*(1), 82-101. doi: 10.1111/j.1467-8489.2011.00565.x

Dugar, Subhasish. (2013). Non- Monetary Incentives and Opportunistic Behavior: Evidence from a Laboratory Public Good Game. *Economic Inquiry, 51*(2), 1374-1388.

Ertan, Arhan, Page, Talbot, & Putterman, Louis. (2009). Who to punish? Individual decisions and majority rule in mitigating the free rider problem. *European Economic Review, 53*(5), 495-511.

Faillo, Marco, Grieco, Daniela, & Zarri, Luca. (2013). Legitimate punishment, feedback, and the enforcement of cooperation. *Games and Economic Behavior, 77*(1), 271-283.

Falk, Armin, Fehr, Ernst, & Fischbacher, Urs. (2005). Driving Forces behind Informal Sanctions. *Econometrica, 73*(6), 2017-2030.

Feeny, David, Berkes, Fikret, McCay, Bonnie J., & Acheson, James M. (1990). The tragedy of the commons: twenty-two years later. *Human ecology, 18*(1), 1-19.

Fehr, Ernst, & Fischbacher, Urs. (2004). Third-party punishment and social norms. *Evolution and Human Behavior, 25*(2), 63-87.

Fehr, Ernst, & Gächter, Simon. (2000). Cooperation and punishment in public goods experiments. *The American Economic Review, 90*(4), 980-994.

Fehr, Ernst, & Gächter, Simon. (2001). Do incentive contracts crowd out voluntary cooperation?

Fehr, Ernst, & Gächter, Simon. (2002). Altruistic punishment in humans. *Nature, 415*(6868), 137-140.

Fehr, Ernst, & Schmidt, Klaus M. (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics, 114*(3), 817-868.

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

Fischbacher, Urs, Gächter, Simon, & Fehr, Ernst. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters, 71*(3), 397-404.

Fisher, A., Wheeler, W. J., & Zwick, R. (1993). Experimental methods in agricultural and resource economics: How useful are they? Agricultural and Resource Economics Review, 22(2), 103-166.

Fonterra. (2007). Fonterra and the environment. Retrieved from https://www.fonterra.com/wps/wcm/connect/81634c804561fdcc8637ff8c696de887/Fonterra+and+the+Environment+A+sustainable+New+Zealand+dairying+industry.pdf?MOD=AJPERES&CACHEID=81634c804561fdcc8637ff8c696de887.

Francisco, Alpízar, Requate, Till, & Schram, Albert. (2004). Collective versus Random Fining: An Experimental Study on Controlling Ambient Pollution. *Environmental and Resource Economics, 29*(2), 231-252.

Frey, Bruno S. (1997). A Constitution for Knaves Crowds out Civic Virtues. *The Economic Journal, 107*(443), 1043-1053.

Fudenberg, Drew, & Pathak, Parag A. (2010). Unobserved punishment supports cooperation. *Journal of Public Economics, 94*(1), 78-86.

Gächter, Simon, Renner, Elke, & Sefton, Martin. (2008). The long-run benefits of punishment. *Science, 322*(5907), 1510-1510.

Gibson, Clark C., Williams, John T., & Ostrom, Elinor. (2005). Local enforcement and better forests. *World Development, 33*(2), 273-284.

Gneezy, Uri, Meier, Stephan, & Rey-Biel, Pedro. (2011). When and why incentives (don't) work to modify behavior. *The Journal of Economic Perspectives, 25*(4), 191-209.

Gordon, H. Scott. (1954). The Economic Theory of a Common-Property Resource: The Fishery. *Journal of Political Economy, 62*(2), 124-142.

Greiner, Ben. (2004). The online recruitment system orsee 2.0-a guide for the organization of experiments in economics. *University of Cologne, Working paper series in economics, 10*(23), 63-104.

Gürerk, Özgür, Irlenbusch, Bernd, & Rockenbach, Bettina. (2006). The competitive advantage of sanctioning institutions. *Science, 312*(5770), 108-111.

Hansen, Lars Gårn. (1998). A damage based tax mechanism for regulation of non-point emissions. *Environmental and Resource Economics, 12*(1), 99-112.

Hardin, Garrett. (1968). The tragedy of the commons. *Science, 162*(3859), 1243-1248.

Harrison, Glenn W. (1988). Predatory pricing in a multiple market experiment: a note. *Journal of Economic Behavior & Organization, 9*(4), 405-417.

Herrmann, Benedikt, Thöni, Christian, & Gächter, Simon. (2008). Antisocial Punishment Across Societies. *Science, 319*(5868), 1362-1367

Hickey, Christopher Wayne, & Martin, Michael Leslie. (2009). *A review of nitrate toxicity to freshwater aquatic species*: Environment Canterbury. Canterbury.

Holmstrom, Bengt. (1982). Moral Hazard in Teams. *The Bell Journal of Economics, 13*(2), 324-340.

Hölmstrom, Bengt. (1979). Moral Hazard and Observability. *The Bell Journal of Economics, 10*(1), 74-91.

Horan, Richard D., Shortle, James S., & Abler, David G. (1998). Ambient taxes when polluters have multiple choices. *Journal of Environmental Economics and Management, 36*(2), 186-199.

Horan, Richard D., Shortle, James S., & Abler, David G. (2002). Ambient Taxes Under m-Dimensional Choice Sets, Heterogeneous Expectations, and Risk-Aversion. *Environmental and Resource Economics, 21*(2), 189-202.

Houser, Daniel, Xiao, Erte, McCabe, Kevin, & Smith, Vernon. (2008). When punishment fails: Research on sanctions, intentions and non-cooperation. *Games and Economic Behavior, 62*(2), 509-532.

Huchet-Bourdon, Marilyne. (2011). Agricultural commodity price volatility.

Hume, David. (1875). *Of the independency of parliament*: Library of Alexandria.

Isaac, R. Mark, McCue, Kenneth F., & Plott, Charles R. (1985). Public Goods Provision in an Experimental Environment. *Journal of Public Economics, 26*(1), 51.

Isaac, R. Mark, & Walker, James M. (1988a). Communication and free- riding behavior: The voluntary contribution mechanism. *Economic Inquiry, 26*(4), 585-608.

Isaac, R. Mark, & Walker, James M. (1988b). Group size effects in public goods provision: The voluntary contributions mechanism. *The Quarterly Journal of Economics, 103*(1), 179.

Isaac, R. Mark, Walker, James M., & Thomas, Susan H. (1984). Divergent evidence on free riding: An experimental examination of possible explanations. *Public Choice (pre-1986), 43*(2), 113.

Janssen, Marco A., Holahan, Robert, Lee, Allen, & Ostrom, Elinor. (2010). Lab experiments for the study of social-ecological systems. *Science, 328*(5978), 613-617.

Karp, Larry. (2005). Nonpoint source pollution taxes and excessive tax burden. *Environmental and Resource Economics, 31*(2), 229-251.

Kingi, Tanira, Park, Simon, & Scarsbrook, Mike. (2012). *Solutions for a Sustainable Lake Rotorua: The Farmers' Perspective*. Paper presented at the New Zealand Landcare Trust Conference.

Kingsley, David C. (2015). Peer punishment across payoff equivalent public good and common pool resource experiments. *Journal of the Economic Science Association, 1*(2), 197-204.

Kube, Sebastian, & Traxler, Christian. (2011). The interaction of legal and social norm enforcement. *Journal of Public Economic Theory, 13*(5), 639-660.

Lazzarini, Sergio G., Miller, Gary J., & Zenger, Todd R. (2004). Order with some law: complementarity versus substitution of formal and informal arrangements. *Journal of Law, Economics, and Organization, 20*(2), 261-298.

Le Grand, Julian. (1997). Knights, knaves or pawns? Human behaviour and social policy. *Journal of social policy, 26*(02), 149-169.

Ledyard, John. (1997). Public goods: A survey of experimental research. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (pp. 111-194). New Jersey: Princeton University Press.

Ledyard, John, & Roberts, John. (1974). On the incentive problem with public goods.

Levitt, Steven D., & List, John A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal about the Real World? Journal of Economic Perspectives, 21(2), 153-174.

Marwell, Gerald, & Ames, Ruth E. (1979). Experiments on the Provision of Public Goods. I. Resources, Interest, Group Size, and the Free-Rider Problem. *American Journal of Sociology, 84*(6), 1335-1360.

Marwell, Gerald, & Ames, Ruth E. (1981). Economists free ride, does anyone else? *Journal of Public Economics, 15*(3), 295-310.

Masclet, David, Noussair, Charles N., & Villeval, Marie-Claire. (2013). Threat and punishment in public good experiments. *Economic Inquiry, 51*(2), 1421-1441.

Matheson, F., Quinn, J. , & Hickey, C. . (2012). Review of the New Zealand instream plant and nutrient guidelines and development of an extended decision-making framework: Phases 1 and 2 final report. Prepared for the Ministry of Science & Innovation Envirolink Fund. *NIWA Client Report*

Meals, Donald W., Dressing, Steven A., & Davenport, Thomas E. (2010). Lag Time in Water Quality Response to Best Management Practices: A Review. *Journal of environmental quality., 39*(1), 85-96.

Meinzen-Dick, Ruth. (2007). Beyond panaceas in water institutions. *Proceedings of the National Academy of Sciences of the United States of America, 104*(39), 15200-15205.

Millock, Katrin, & Salanie, Francois. (2005). Nonpoint Source Pollution when Polluters Might Cooperate. *B.E. Journal of Economic Analysis and Policy: Topics in Economic Analysis and Policy, 5*(1), 1-25.

Ministry for the Environment. (2015a). *A Guide to the National Policy Statement for Freshwater Management 2014*. Wellington: Ministry for the Environment.

Ministry for the Environment. (2015b). *Making Collaborative Groups Work: A guide for those involved in collaborative processes*. Wellington: Ministry for the Environment.

Monaghan, R. M., Wilcock, R. J., Smith, L. C., Tikkisetty, B., Thorrold, B. S., & Costall, D. (2007). Linkages between land management activities and

water quality in an intensively farmed catchment in southern New Zealand. *Agriculture, Ecosystems & Environment, 118*(1–4), 211-222.

Murphy, Ryan O., Rapoport, Amnon, & Parco, James E. (2006). The breakdown of cooperation in iterative real-time trust dilemmas. *Experimental Economics, 9*(2), 147-166.

Musgrave, R. A. (1959). *The theory of public finance*. New York: McGraw-Hill.

Nash, John. (1951). Non-cooperative games. *Annals of mathematics*, 286-295.

New Zealand Government. (2014). National Policy Statement for Freshwater Management 2014.

New Zealand Landcare Trust. (2012). Aorere Catchment. Retrieved from http://www.landcare.org.nz/Regional-Focus/Nelson-Office/Aorere-Catchment

Nikiforakis, Nikos, & Normann, Hans-Theo. (2008). A comparative statics analysis of punishment in public-good experiments. *Experimental Economics, 11*(4), 358-369.

Nikiforakis, Nikos, Normann, Hans-Theo, & Wallace, Brian. (2010). Asymmetric enforcement of cooperation in a social dilemma. *Southern Economic Journal, 76*(3), 638-659.

Nikiforakis, Nikos, Noussair, Charles N., & Wilkening, Tom. (2012). Normative conflict and feuds: The limits of self-enforcement. *Journal of Public Economics, 96*(9), 797-807.

Noussair, Charles, & Tucker, Steven. (2005). Combining monetary and social sanctions to promote cooperation. *Economic Inquiry, 43*(3), 649-660.

Organisation for Economic Co-operation and Development. (2007). Instrument Mixes Addressing Non-point Sources of Water Pollution. Retrieved from http://www.oecd.org/env/tools-evaluation/instrumentmixesaddressingnon-pointsourcesofwaterpollution.htm.

Organisation for Economic Co-operation and Development. (2012). Water quality and agriculture: meeting the policy challenge. Paris: OECD.

Olson, Mancur. (1965). The logic of collective action; public goods and the theory of groups.

Organization, World Health. (2004). *Guidelines for drinking-water quality: recommendations* (Vol. 1): World Health Organization.

Ostrom, Elinor. (1994). *Institutional analysis, design principles and threats to sustainable community governance and management of commons.* Paper presented at the ICLARM Conf. Proc.

Ostrom, Elinor. (1999). Coping with tragedies of the commons. *Annual Review of Political Science, 2*(1), 493.

Ostrom, Elinor. (2006). The value-added of laboratory experiments for the study of institutions and common-pool resources. *Journal of Economic Behavior & Organization, 61*(2), 149-163.

Ostrom, Elinor, Burger, Joanna, Field, Christopher B., Norgaard, Richard B., & Policansky, David. (1999). Revisiting the commons: Local lessons, global challenges. *Science, 284*(5412), 278-282.

Ostrom, Elinor, Walker, James, & Gardner, Roy. (1992). Covenants with and without a Sword: Self-governance Is Possible. *American Political Science Review, 86*(02), 404-417.

Page, Talbot, Putterman, Louis, & Unel, Bulent. (2005). Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency*. *The Economic Journal, 115*(506), 1032-1053.

Pannell, David J. (2006). Flat Earth Economics: The Far-reaching Consequences of Flat Payoff Functions in Economic Decision Making. *Review of Agricultural Economics, 28*(4), 553–566.

Parfitt, R. L., Stevenson, B. A., Dymond, J. R., Schipper, L. A., Baisden, W. T., & Ballantine, D. J. (2012). Nitrogen inputs and outputs for New Zealand from 1990 to 2010 at national and regional scales. *New Zealand Journal of Agricultural Research, 55*(3), 241-262.

Parliamentary Comissioner for the Environment. (2012). *Water quality in New Zealand: Understanding the science*. Wellington: Parliamentary Comissioner for the Environment.

Poe, Gregory L., Schulze, William D., Segerson, Kathleen, Suter, Jordan F., & Vossler, Christian A. (2004). Exploring the performance of ambient-based policy instruments when nonpoint source polluters can cooperate. *American Journal of Agricultural Economics, 86*(5), 1203-1210.

Posner, Eric A. (1996). Law, Economics, and Inefficient Norms. *University of Pennsylvania Law Review, 144*(5), 1697-1744.

Posner, Richard A. (1997). Social Norms and the Law: An Economic Approach. *The American Economic Review, 87*(2), 365-369.

Przepiorka, Wojtek, & Diekmann, Andreas. (2013). Individual heterogeneity and costly punishment: a volunteer's dilemma. *Proceedings of the Royal Society B: Biological Sciences, 280*(1759).

Putterman, Louis, Tyran, Jean-Robert, & Kamei, Kenju. (2011). Public goods and voting on formal sanction schemes. *Journal of Public Economics, 95*(9), 1213-1222.

Reuben, Ernesto, & Riedl, Arno. (2013). Enforcement of contribution norms in public good games with heterogeneous populations. *Games and Economic Behavior, 77*(1), 122-137.

Romstad, Eirik. (2003). Team approaches in reducing nonpoint source pollution. *Ecological economics, 47*(1), 71-78.

Samuelson, Paul A. (1954). The Pure Theory of Public Expenditure. *The Review of Economics and Statistics, 36*(4), 387-389.

Seabright, Paul. (1993). Managing Local Commons: Theoretical Issues in Incentive Design. *The Journal of Economic Perspectives, 7*(4), 113-134.

Segerson, K. (1988). Uncertainty and incentives for nonpoint pollution control. *Journal of Environmental Economics and Management, 15*(1), 87-98.

Segerson, Kathleen, & Wu, JunJie. (2006). Nonpoint pollution control: Inducing first-best outcomes through the use of threats. *Journal of Environmental Economics and Management, 51*(2), 165-184.

Shinada, Mizuho, & Yamagishi, Toshio. (2007). Punishing free riders: direct and indirect promotion of cooperation. *Evolution and Human Behavior, 28*(5), 330-339.

Shinada, Mizuho, Yamagishi, Toshio, & Ohmura, Yu. (2004). False friends are worse than bitter enemies: “Altruistic” punishment of in-group members. *Evolution and Human Behavior, 25*(6), 379-393.

Shortle, James S., & Horan, Richard D. (2001). The economics of nonpoint pollution control. *Journal of Economic Surveys, 15*(3), 255-289.

Spraggon, John. (2002). Exogenous targeting instruments as a solution to group moral hazards. *Journal of Public Economics, 84*(3), 427-456.

Suter, Jordan F., Vossler, Christian A., & Poe, Gregory L. (2009). Ambient-based pollution mechanisms: A comparison of homogeneous and heterogeneous groups of emitters. *Ecological economics, 68*(6), 1883-1892.

Suter, Jordan F., Vossler, Christian A., Poe, Gregory L., & Segerson, Kathleen. (2008). Experiments on damage-based ambient taxes for nonpoint source polluters. *American Journal of Agricultural Economics, 90*(1), 86-102.

Synlait. (2016). Lead with pride. Retrieved from http://www.oecd.org/env/tools-evaluation/instrumentmixesaddressingnon-pointsourcesofwaterpollution.htm.

Traulsen, Arne, Röhl, Torsten, & Milinski, Manfred. (2012). An economic experiment reveals that humans prefer pool punishment to maintain the commons. *Proceedings of the Royal Society B: Biological Sciences, 279*(1743), 3716-3721.

Van de Kragt, Alphons J. C., Dawes, Robyn M., & Orbell, John M. (1988). Are People Who Cooperate 'Rational Altruists'? *Public Choice, 56*(3), 233-247.

Van Soest, D. P., & Vyrastekova, J. (2006). Peer enforcement in CPR experiments: The relative effectiveness of sanctions and transfer rewards, and the role of behavioral types. *Using Experimental Methods in Environmental and Resource Economics*, 113-136.

Varughese, George, & Ostrom, Elinor. (2001). The contested role of heterogeneity in collective action: some evidence from community forestry in Nepal. *World Development, 29*(5), 747-765.

Vedeld, T. (2000). Village politics: heterogeneity, leadership and collective action. *Journal of Development Studies, 36*(5), 105-134.

Vossler, Christian A., Poe, Gregory L., Schulze, William D., & Segerson, Kathleen. (2006). Communication and incentive mechanisms based on group performance: an experimental study of nonpoint pollution control. *Economic Inquiry, 44*(4), 599-613.

Walker, J. M., Gardner, Roy, Herr, Andrew, & Ostrom, Elinor. (2000). Collective Choice in the Commons: Experimental Results on Proposed Allocation Rules and Votes. *The Economic Journal, 110*(460), 212-234.

Wilson, David Sloan, Ostrom, Elinor, & Cox, Michael E. (2013). Generalizing the core design principles for the efficacy of groups. *Journal of Economic Behavior & Organization, 90*, S21-S32.

Xepapadeas, A. P. (1992). Environmental policy design and dynamic nonpoint-source pollution. *Journal of Environmental Economics and Management, 23*(1), 22-39.

Xepapadeas, A. P. (1995). Observability and choice of instrument mix in the control of externalities. *Journal of Public Economics, 56*(3), 485-498.

Yamagishi, Toshio. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and social Psychology, 51*(1), 110-116.

Zhang, Boyu, Li, Cong, De Silva, Hannelore, Bednarik, Peter, & Sigmund, Karl. (2014). The evolution of sanctioning institutions: an experimental approach to the social contract. *Experimental Economics, 17*(2), 285-303.

8. Appendix: Instructions for a 1:9 punishment experiment

Instructions

This is an experiment on decision making. If you read the following instructions carefully, you can, depending on your decisions, earn a considerable amount of money. All earnings on your computer screens are in Experimental Dollars. These Experimental Dollars will be converted to real Dollars at the end of the experiment, at a rate of

___ Experimental Dollars = 1 real Dollar.

Today's session will be conducted using the computer network located here in this laboratory. It will be divided into 10 different periods. Attached to these instructions you will find a sheet labelled Personal Record Sheet, which will help you keep track of your earnings based on the decisions you might make. You are not to reveal this information to anyone. It is your own private information.

You have been assigned to a group of four (yourself and three other) participants. This will be your group for the entire session.

The First Stage each Period

At the beginning of each period each participant receives an endowment of 12 tokens. In **Stage 1** each period you (and the others in your group) must decide how many tokens to place into either or both of 2 accounts: a private account and a group account. All tokens must be placed in one account or the other. Each token you place in the private account generates a return to you (independent of what anyone else does), and each token you place in the group account generates a return that depends on how many tokens that others in your group place in the group account. Your earnings in a period are the sum of your earnings from the

private account and your earnings from the group account, minus 12. Returns to the two accounts are listed on your input screen as shown on the next page. Everybody has the same returns.

You and all the other members of your group will each get a share of the total group earnings that depends on your token placements. If a total of X tokens are placed in the group account by group members, then the total group payoff is $18X - 0.4X^2$ Experimental Dollars. For example, if a total of 10 tokens are placed in the group account, the total group payoff is $(18\times10) - (0.4\times100) = 180 - 40 = 140$. This amount, along with every other possible total token placement in the group account is shown on your input screen below.

Your share of this total group payoff equals your number of tokens placed in the group account as a fraction of the total tokens you and the others in your group place in the group account. For example, if a total of 10 tokens are placed in the group account, and you placed 2 of these 10 tokens, then you receive $2/10 = 0.2$ of the 140 total group payoff, or $0.2 \times 140 = 28$.

Your private account generates a return to you that depends only on your tokens placed in the private account. In particular, you will receive 2 Experimental Dollars for each token that you place in your private account.

You will indicate your decisions on the input-screen for the first stage:

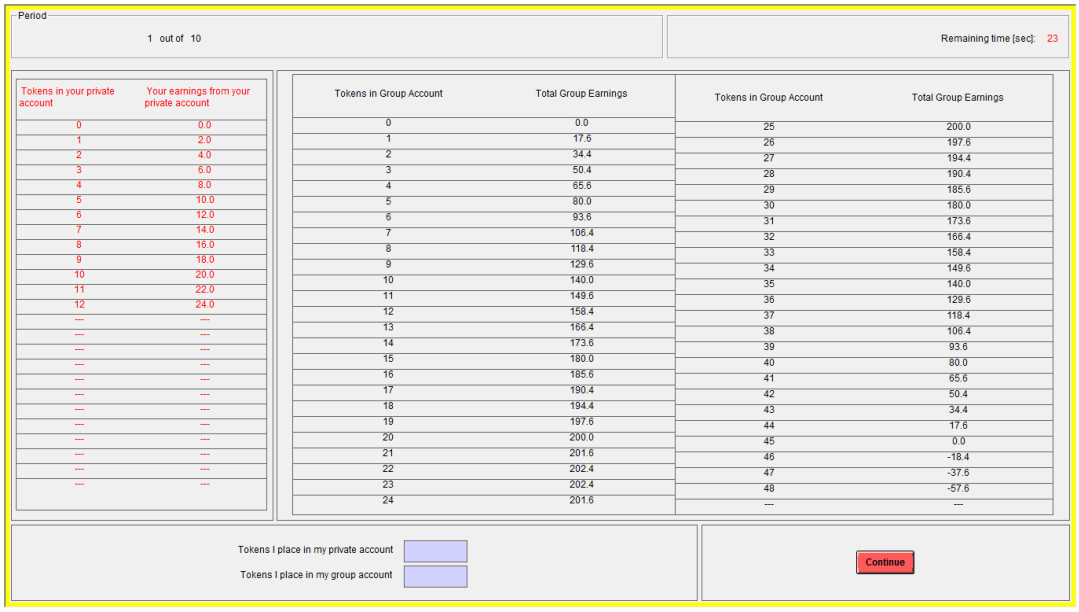

Since your endowment each period is 12 tokens, the two numbers you indicate on your input-screen must be whole numbers between 0 and 12 and must sum to 12. After entering your decision you must press the **Continue** button. Once you have done this, your decision has been made and cannot be changed.

After all participants in your group have made their decisions the following income screen will show you the total amount of tokens placed in the group account by all four participants of your group (including you). Also this screen shows your earnings for the first stage of the period. Your earnings are the sum of your earnings from the private account and your earnings from the group account. During the experiment you will record this information on your hardcopy record sheet and then click the **Continue** button.

Detail of Results Screen for First Stage

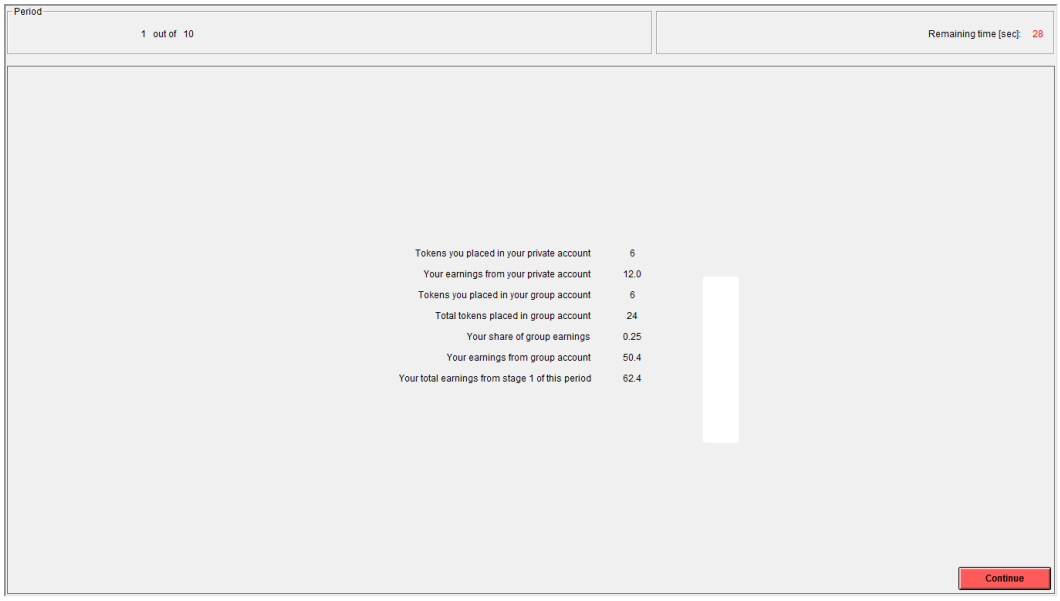

The Second Stage each Period

In the second stage you will see the amount of tokens placed in the group account by all four participants of your group. Moreover, in this stage you can decide whether to **decrease** the earnings received from the first stage of the others in your group by assigning **deduction points**. These other participants can also decrease your first stage earnings if they wish to. This is apparent from the input screen at the second stage, shown below.

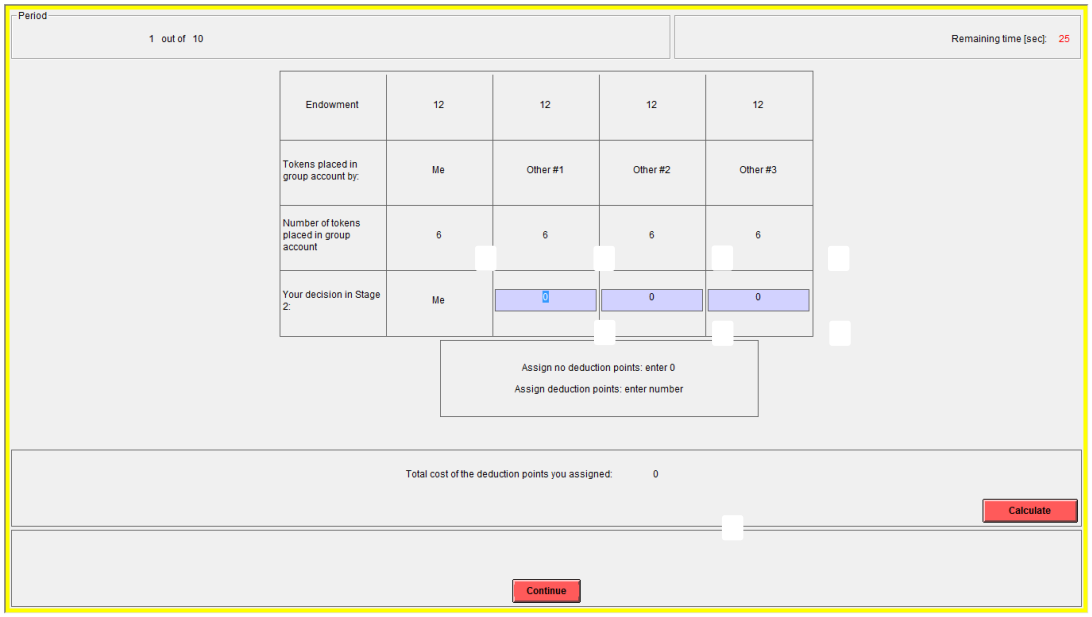

Your allocation to the group account is displayed **in the first column**, while the allocations to the group account **by the other people** are shown in the **remaining three columns**. Note that the order in which others' allocations are displayed will be determined at random in every period. The allocation in the second column, for example, could represent a different person in different periods. The same holds true for the other two columns.

You will have to decide how many deduction points to assign to **each** of these other three participants in your group. You must enter a number for each of them. If you do not wish to change the earnings of a person in your group then you must enter 0. You can **assign up to 5 points to each participant**.

You will incur costs from assigning deduction points. Every deduction point you assign costs you 1 Experimental Dollar. For example, if you assign 2 deduction points to one person, this costs you 2 Experimental Dollars; if, in addition, you assign 4 deduction points to another person this costs you an additional 4 Experimental Dollars. In total for this example you will have assigned 6 points and your **total costs** therefore amount to 6 Experimental Dollars.

After you have assigned points to each of the other three participants you can click the button "**Calculate**" (see the second stage input screen). On the screen you will then see the total costs of your assigned points. As long as you have not yet clicked the **Continue** button, you can still change your decision. To recalculate the costs after a change of your assigned points, simply press the "Calculate" button again.

If you assign 0 deduction points to a particular participant (i.e., enter "0"), you will not alter his or her first stage earnings. However, if you assign **one deduction point** to a participant you will **decrease** his or her first stage earnings by **9 Experimental Dollars**. If you assign a participant **2 deduction points** you will **decrease** his or her first stage earnings by **18 Experimental Dollars**, and so on. Each deduction point that you assign to another person will reduce his or her first stage earnings by 9 Experimental Dollars. Similarly, each deduction point assigned to you by another participant will reduce your first stage earnings by 9 Experimental Dollars:

Costs of received deduction points = 9 × Sum of received deduction points.

How much the first stage earnings are decreased depends on the sum of deduction points received. For instance, if somebody receives **a total** of **3 deduction points** (from all other participants in this period), his or her earnings would be decreased by **27 Experimental Dollars**. If somebody receives a total of **4 deduction points,** his or her earnings are reduced by **36 Experimental Dollars**. Your total earnings from the two stages are therefore calculated as follows:

Total earnings at the end of the second stage = period earnings =
= Earnings from the first stage – 9 × (sum of received deduction points)
– (sum of deduction points you have assigned) – 12

Note that everyone has a fixed amount of 12 Experimental Dollars subtracted every period.

After all participants have made their decision, your earnings from the period will be displayed on a screen such as the one shown below. After you have viewed the earnings screen the period is over and the next period commences. Recall that 10 periods will be conducted.

Recording Rules

During every period you should write down the information shown on your results screens on your Personal Record Sheet. The Stage 1 results screen shows the group account allocations you should record, and the final earnings screen shows your Deduction Points and period earnings. Be sure to record your total earnings for each period in the rightmost column.

Summary

1. All subjects use the same Earnings Tables.
2. In each period, you and every other participant will each have 12 tokens to allocate.
3. In each period, you will decide how many tokens to place in your private account and how many to place in your group account. You must allocate all 12 tokens each period.
4. Your earnings from the private account depend only on your decision about how many tokens to place in this account.
5. Your earnings from the group account depend upon how many tokens you and the other three participants of your group place in this account. You receive a share of group earnings that depends the fraction of the total tokens in the group account that you placed in the group account.
6. You may assign up to 5 Deduction Points to each of the other individuals in your group. Each point you assign costs you 1 Experimental Dollar.
7. These other individuals can assign Deduction Points to you. Each Point assigned to you reduces your earnings by 9 Experimental Dollars.
8. You will interact with the same three other individuals for all decision periods.
9. Results and earnings should be recorded on your Record Sheet at the end of each period.

Detail of Earnings Screen at the end of the Second Stage:

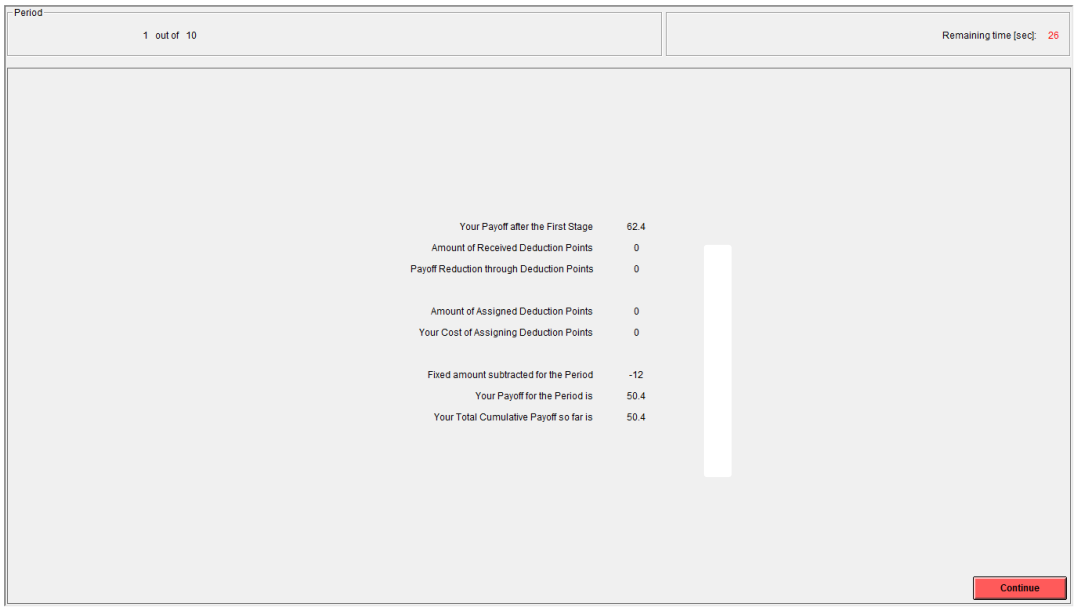

Personal Record Sheet

Period number	My Tokens in Private Account	My Earnings from Private Account	My Tokens in Group Account	Total Tokens in Group Account	My Share of Group Earnings	My Earnings from Group Account	My Total Earnings from Stage 1	Amount of Received Deduction Points	Payoff Reduction through Deduction Points	Cost of Assigning Deduction Points	Payoff for the Period	Cumulative Payoff so far
1												
2												
3												
4												
5												
6												
7												
8												
9												
10												