INFORMATION IN GAMES

by

Joshua Scott Cherry

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2011

Doctoral Committee:

Assistant Professor Daisuke Nakajima, Co-Chair
Professor Lones A. Smith, Co-Chair, University of Wisconsin
Professor Tilman M. Borgers
Associate Professor Uday Rajan
Assistant Professor Yusufcan Masatlioglu

to my parents
for always supporting my pursuit of education

ACKNOWLEDGEMENTS

I first wish to thank my committee — Tilman Borgers, Yusufcan Masatlioglu, Daisuke Nakajima, Uday Rajan, and especially Lones Smith — for being so generous with their time and advice during my experience in graduate school. The first two chapters of my dissertation are coauthored with Lones Smith. Through our collaboration he has taught me not only a great deal of economics, but more importantly, how to conduct research and communicate results.

The third chapter is coauthored with Steve Salant and Neslihan Uler. I truly appreciated the opportunity to work with both; they are great scholars, and even better people. I would also like to thank Logan Lewis, Dmitry Lubensky, Todd Pugatch, Collin Raymond, Doug Smith, and Mike Stevens for being wonderful friends and colleagues.

I am indebted to my undergraduate advisors Steve Goldman and Botond Koszegi for preparing me for graduate school and fostering in me a passion for the discipline. I would also like to thank the National Science Foundation, Rackham Graduate School, and Lones Smith for their generous financial support. My work has benefited from the Incentive Centered Design lab group, and in particular Jeff Mackie-Mason. It has also been improved by seminar participants at the University of Michigan, Midwest Economic Theory Conference, NSF/NBER/CEME Conference and elsewhere.

I am especially thankful for the care, encouragement, and emotional support I received from Laura Kawano. Finally, I most grateful to my family for their unwavering confidence in me and enduring willingness to help me pursue my goals.

TABLE OF CONTENTS

DEDICATION . . . . **ii**

ACKNOWLEDGEMENTS . . . . **iii**

LIST OF FIGURES . . . . **vi**

LIST OF TABLES . . . . **vii**

CHAPTER

I. Unattainable Payoffs for Repeated Games of Private Monitoring 1

1.1 Introduction 1
1.2 Motivational Example 4
1.3 A Mediated Repeated Game 7
1.4 Markov Perfect Correlated Equilibrium 8
1.5 Repeated Games of Private Monitoring 12
1.6 Unattainable Private Monitoring Payoffs 14
1.7 Conclusion 17

II. Strategically Valuable Information 18

2.1 Introduction 18
2.2 Model 19
2.2.1 The Informational Setting 19
2.2.2 Formulation as Signals 22
2.2.3 The Strategic Environment 23

2.3 The Equivalence Result . . . 23
2.4 Economic Applications . . . 24
2.4.1 Application to Repeated Games . . . 24
2.4.2 Sunspots in General Equilibrium . . . 26
2.5 Conclusion . . . 26

III. Electing Efficiency . . . 28

3.1 Introduction . . . 28
3.2 The Experiment . . . 31
3.2.1 Design and Procedures . . . 31
3.3 Theoretical Predictions and Hypotheses . . . 34
3.4 Data Analysis . . . 36
3.4.1 Exogenous Groups and Investment Decisions . . . 36
3.4.2 Voting for Group Size: The Plurality Rule . . . 39
3.5 Discussion and Conclusions . . . 43

IV. Omitted Proofs . . . 45

4.1 Any PCE Payoff is an MPCE Payoff: Proof of Lemma I.1 . . . 45
4.2 Characterization of MPCE: Proof of Theorem I.1 . . . 46
4.3 Algorithm: Proof of Theorem I.2 . . . 47
4.4 MPCE as an Upper Bound: Proof of Theorem I.3 . . . 48
4.5 MPCE Inclusion is Tight: Proof of Theorem I.4 . . . 50
4.6 Sufficiency . . . 50
4.7 Necessity . . . 50

Bibliography . . . **52**

LIST OF FIGURES

Figure

1.1	Example Stage Game	4
1.2	Payoffs in the Repeated Game in Figure 1.1.	11
1.3	Prisoner's Dilemma ($b > 1, c > 0$)	16
2.1	Information Structure $\mathcal{P}$	20
2.2	Information Structure $\mathcal{Q}$	20
2.3	Sunpot with Two Ouctomes	21
2.4	Sunspot with Four Outcomes	21
2.5	Information Structure $\mathcal{P}$	21
2.6	Information Structure $\mathcal{Q}$	21
2.7	Information Structure $\mathcal{R}$ is constructed by embedding $\mathcal{Q}$ in both halves of $\mathcal{P}$	21
2.8	Pure Coordination Game	24
2.9	Game of Chicken	24
3.1	Situation Analyzer for Groups of Two	34
3.2	Mean Investment	37

LIST OF TABLES

Table

3.1 Experimental Design 33
3.2 Predicted Mean Investment in Project B 35
3.3 Predicted versus Observed Mean Investment 38
3.4 Predicted versus Observed Mean Payoff 38
3.5 Ordinary Least Squares Results 40
3.6 Percentage of Votes in Part V 41
3.7 Multinomial Logit Regression – Marginal Effects 41
3.8 Mean Investment and Payoff Conditional on Chosen Group Size 42
3.9 A Comparison of Efficiency Levels 43

CHAPTER I

Unattainable Payoffs for Repeated Games of Private Monitoring

1.1 Introduction

A repeated game is a stylized model of a long-term relationship. The most common solution concepts for repeated games are Subgame Perfect Equilibrium (SPE) and its extension to imperfect public monitoring, Perfect Public Equilibrium (PPE). In both cases, equilibrium strategies depend only on commonly observed histories. This yields a recursive property that *every continuation game is equivalent to the entire game*. Abreu, Pearce, and Stacchetti (APS) pursued this logic in 1986 and 1990, and so characterized equilibrium payoffs using methods inspired by dynamic programming.

APS built on Green and Porter's 1984 seminal exposition of dynamic Cournot oligopoly — who in turn took inspiration from Stigler's influential (1964) theory of dynamic Bertrand oligopoly. To sustain collusion in a world with hidden pricing, Stigler had proposed that firms initiate a price war if standard statistical tests suggested cartel cheating. Stigler struggled with a problem that afflicts much of economic theory — i.e, *any* dynamic setting with unobserved payoff-relevant actions that do not just feed into an observable stochastic aggregate, like an observed price. It is arguably important in all long-term partnerships ranging from relational contracting to international political relations. Restricting attention to public signals intuitively ignores strategically relevant information, and misses the potential richness of the dynamic structure. Upon reflection, public monitoring should only be justified as a tractable *approximation* of this richer private monitoring "reality". So then, how restrictive is it? Exactly how much does private monitoring expand the scope for collusion in oligopoly, eg? Our finding is: substantially, in some cases.

Unfortunately, not only is private monitoring an interesting problem, it is also a difficult one. And thus Stigler's question remains unsolved after so much time. For as is well-known, private monitoring in repeated games induces correlated private histories, and this frustrates attempts to use recursive methods, as in APS. On a sequential equilibrium path, continuation play in any period constitutes a correlated equilibrium, where the private histories act as endogenous correlation devices. And computing a best response in a non-trivial sequential equilibrium may well require an impossibly complicated probabilistic inference. It is no surprise that Kandori (2002) calls this "one of the best known long-standing open questions in economic theory."

Though incentives are harder to provide with unobserved actions, the induced correlation may facilitate coordination (as in Aumann 1974, 1987), and augment the sequential equilibrium payoff set. So motivated, we explicitly incorporate correlated private histories, as first studied by Lehrer (1992). But our approach admits arbitrary correlation each period. First, we develop a new solution concept for infinitely repeated games with perfect monitoring that reflects these correlation possibilities. Whereas APS defined an operator that took the Nash equilibria of the 'auxiliary game' at the start of every subgame, we take correlated equilibria. This is a well-defined exercise since we publicize past correlated recommendations. The largest fixed point of the resulting operator yields the *Markov Perfect Correlated Equilibrium* (MPCE) payoff set, and is therefore recursive like PPE. Notably, not only is our solution concept tractable, it is arguably easier to compute than is the PPE set. For unlike Nash equilibrium, the set of correlated equilibria can be found by linear methods.

We then explore the implications of MPCE for repeated games of private monitoring. We show that for any monitoring structure, the set of sequential equilibrium payoffs is contained within the MPCE payoff set for the corresponding *expected stage game*. This helps us deduce the tightest bound on repeated game equilibrium payoffs that is independent of the monitoring structure.

Our paper has two parts. We begin with an infinitely repeated game of observed actions, and embellish it with an extensive-form correlation device that can generate any (possibly) history-dependent private messages every period. Since messages are made public after players act, a recursive structure emerges. Unlike Prokopovych (2006) who first took this road, we then show that a Markovian device suffices to describe all equilibrium payoffs. This yields our MPCE solution concept. Theorem I.1 characterizes the resulting payoff set — it is compact, convex, and nondecreasing in the discount factor. Also, it contains all subgame perfect payoffs. Theorem I.2 describes a tractable, recursive algorithm for computing it.

In the second thrust, we turn to a repeated game of private monitoring, and relate its sequential equilibria to the MPCE of the corresponding repetition of the expected stage game. Theorem I.3 asserts that this set serves as an upper bound for the sequential equilibrium payoffs. We thereby identify the certainly unattainable sequential equilibrium payoffs for a repeated game of private monitoring for any fixed discount factor. Notably this bound holds for *all* monitoring structures, as well as private strategies in public monitoring games. In other words, we precisely compute the set of payoffs potentially added by the richer information structure introduced by private monitoring — one possible completion of Stigler's original thought.

Theorem I.4 explores how our payoff upper bound can be made tight. For unlike MPCE, a standard repeated game of private monitoring with an initial period does not allow any pre-play signals. So motivated, we augment the MPCE concept. We compute the Nash equilibrium payoffs of all auxiliary games using continuation payoffs drawn from the MPCE set. Put differently, this applies the APS operator to the MPCE payoff set. Any payoff in the resulting set can be supported as a sequential equilibrium in a repeated game with some monitoring structure. We therefore obtain the *tightest possible bound* that makes no reference to the monitoring structure.

Research on repeated games with private monitoring has been driven by the folk theorem, and so proceeded by finding computable classes of sequential equilibria. In contrast, we provide a *superset* of the equilibrium payoff set. The earliest work found nearly efficient equilibria that dispense with all but a simple summary of past play. Loosely, these "belief-based" approaches focus on the chance of misleading private messages. This is possible when the monitoring is sufficiently accurate (e.g. Sekiguchi 1997, and Bhaskar and Obara 2002). A clever and recursive set of non-trivial equilibria in which players' beliefs are irrelevant was later identified by Piccione (2002) and Ely and Valimaki (2002), and greatly extended by Ely, Horner, and Olszewski (2005). While this belief-free approach constitutes a strict subset of all sequential equilibrium payoffs and requires sufficiently patient players, it often secures a folk theorem.

Our paper is not intended *in any way* as a contribution to the folk theorem literature. For we shift from characterizing what is a sequential equilibrium, to what is *not*. Abreu, Milgrom, and Pearce (1991) call into question the relevance of a folk theorem in this setting. Since a discounted repeated game unjustifiably entwines time preference and the frequency of monitoring, the discrete time folk theorem logic yields more informative monitoring with more rapid play. A large discount factor is an appropriate modeling choice only if opportunities to observe others' actions are frequent. Though Coca Cola and Pepsi can change prices arbitrarily often, without similarly (and implausibly) frequent reports of their rivals' actions, they will change behavior only as often as information arrives. The analysis of dynamic oligopoly in Green and Porter (1984) was meaningful precisely because of the fixed discount factor. Our analysis sheds light on equilibrium payoffs when the folk theorem does not apply — such as when interaction is not very frequent, or when information revelation about unobserved actions inherently cannot be accelerated. Instead our paper offers definitive insights on payoffs for those applications with a fixed discount factor.

In arguing that the Cournot-Nash outcome was the wrong benchmark for deducing collusion Porter (1983) wrote: "Industrial organization economists have recognized for some time that the problem of distinguishing empirically between collusive and noncooperative behavior, in the absence of a 'smoking gun', is a difficult one." Firms can achieve higher payoffs in a fully compliant, noncooperative fashion. Combining this insight with our approach, we allow that firms might avail themselves of correlated information, and potentially achieve more outcomes. Our MPCE solution concept is agnostic about the details of who knows what and when. In this way, MPCE is a better litmus test of cheating for regulators to rule out the possibility of collusion; otherwise, one might mistakenly assert an antitrust violation.

The paper is organized as follows. We gently begin with a motivational example. Next, we discuss infinitely repeated games of perfect monitoring with an extensive form correlation device, and develop our new MPCE solution concept. We illustrate it by returning to our example. We then formally describe infinitely repeated games with private monitoring, and compare their sequential equilibrium payoffs with the MPCE payoffs of standard repeated games. Here, we establish our payoff upper bound and show that it can be tight. All proofs are in the Appendix.

1.2 Motivational Example

A. Analysis of a Repeated Prisoners' Dilemma.

Consider an infinitely repeated two player game of perfect monitoring with payoffs given by Figure 1.1. The players share the discount factor 3/4, and so are not patient enough to support the cooperative outcome in a subgame perfect equilibrium. Stahl (1991) shows that even with public correlation, the set of SPE payoffs is the convex hull of $\{(0,0),(7,0),(0,7)\}$, and thus the highest symmetric subgame perfect equilibrium payoff is $(7/2, 7/2)$. If instead we have imperfect public monitoring, then from Kandori (1992) the PPE payoff set is even smaller.[1]

	C	D
C	(4,4)	(-13,20)
D	(20,-13)	(0,0)

Figure 1.1: Example Stage Game

[1] Kandori (1992) shows that the PPE set is monotone in the informativeness (in the sense of Blackwell (1953)) of the public signal. Specifically, the PPE payoff set weakly shrinks when the public signal is garbled.

Next, suppose that players privately observe a payoff irrelevant signal from $\{g, b\}$ before play each period. The signal profiles $\{(g, g), (g, b), (b, g)\}$ occur with probabilities $(1/2, 1/4, 1/4)$, respectively, independently of the past. After actions are chosen, the private signal profile is commonly revealed to both players. To simplify matters, assume players can access a public correlation device that draws a number z from a uniform distribution on $[0, 1]$.

Consider the strategy profile: "In phase 1, play C after observing g, and D after b. If agents play the same action, then repeat phase 1. Otherwise, if player $i = 1, 2$ alone plays D, then proceed to phase 2-i, where player i plays C and player $-i$ mixes so that i gets an expected payoff of 0. If both players play C, then stay in phase 2-i. Otherwise, return to phase 1."

When the repeated game is enriched by the signal process, these strategies constitute a sequential equilibrium. The equilibrium payoff for each player is

$$v = (1/4)(4(1/2) - 13(1/4) + 20(1/4)) + (3/4)(v(1/2) + 2v(1/4) + 0(1/4))$$

i.e. $v = 15/4$. When called upon to play C, a player will acquiesce because

$$(1/4)(4(1/2) - 13(1/4)) + (3/4)((15/4)(1/2) + 2(15/4)(1/4)) \geq (1/4)20(1/2)$$

At the start of phase 1, both players expect the payoff $15/4$. In phase 2-i, player i expects a payoff of 0 and player $-i$ expects $15/2$. The payoff $(15/4, 15/4)$ Pareto dominates the highest symmetric subgame perfect equilibrium payoff $(7/2, 7/2)$ attainable without any signals. In fact, $(15/4, 15/4)$ dominates any symmetric PPE payoff attainable under any imperfect public monitoring structure. Nevertheless, $(15/4, 15/4)$ can be attained in an MPCE because both the information and strategies depend only on the most recent period.

This example reflects two truths: (a) relative to public monitoring, private monitoring may greatly expand the set of sequential equilibrium payoffs, and (b) MPCE captures these richer information structures and the larger payoff set.

For a bigger picture insight, consider the intuition in Kandori (2008). Although correlation cannot enhance play in the one-shot prisoners' dilemma, the repeated game instead confronts players with a game of chicken. This auxiliary game admits nontrivial correlated equilibria. Thus, imperfectly correlated signals can have a meaningful dynamic strategic effect.

More specifically, in this game the gain to defecting is higher when the other player cooperates than when he defects since $20 - 4 = 16 > 0 - (-13) = 13$. But our correlating signal confuses the players about what action profile is played in any period. Consequently, the temptation to cheat is a weighted average of 16 and 13, and so smaller than if no correlation were available. This correlation is not without a cost, since the equilibrium prescribes the most efficient payoff $(4, 4)$ less often.

B. Economic Settings

We now argue that this example captures a wide range of economic settings.

Repeated Partnership. A theorist and an empiricist seek to write a paper together. At the start of each day, they independently choose whether to exert high effort or low effort (actions C and D in the example,respectively). They meet at the end of every day to demonstrate their accomplishments. Suppose, however, that they entertain subjective interpretations of their colleague's effort (as in MacLeod (2003) and Fuchs (2007)). Each colleague entertains either a good (g) or bad (b) subjective interpretation, corresponding to high or low effort by his colleague, respectively. For example, the empiricist's regression output is commonly observed, but the theorist cannot accurately gauge the effort required to produce the results. A key additional source of discounting here is that the partnership might end.

Principal-Agency. An employee chooses each period to exert either high or low effort (actions C and D in the example, respectively). His manager simultaneously chooses one of two compensation schemes: pay a bonus for high output, or never pay a bonus (actions C and D in the example). The private signals can take one of the following two interpretations. In the first, private signals are non-binding recommendations to managers and employees made by a board of directors. The board's fiduciary duty to maximize shareholder value would justify influencing the relational contracts implemented by the firm. In the second, the private signals are subjective evaluations of output made by the agents. MacLeod (2003) characterizes the optimal contract when the joint density of the subjective evaluations is given. With MPCE one can study this context while being agnostic about the exact structure of the agents' subjectivity.

Dynamic Quality Choice. A single product firm has one long-run customer and can use higher or lower quality inputs (actions C and D). A product with better inputs yields higher performance. Without observing the firm's input choice, the customer decides whether or not to purchase the item (actions C and D). After each period, the firm and the customer each observe a private signal indicating a good (g) or bad (b) performing product.

Secret Price Cuts. Thus the actions C and D in the inspirational example from Stigler (1964) represent high and low prices, while the private signals g and b may correspond to high and low demand.

1.3 A Mediated Repeated Game

We begin with a repeated game of perfect monitoring $G(\delta)$, played in periods $1, 2, \ldots$, and payoffs discounted by the factor $0 < \delta < 1$. Each period, every player $i \in N = \{1, 2, \ldots, n\}$ chooses an action a_i from a finite action set A_i. An action profile a is thus an element of $A = \prod_i A_i$, the set of pure action profiles.[2] Payoffs given the action profile a are $u(a) = (u_1(a), \ldots, u_n(a))$. Let α_i denote the mixed action for player i that chooses action $a_i \in A_i$ with chance $\alpha_i(a_i)$. Abusing notation, $u(\alpha) = (u_1(\alpha), \ldots, u_n(\alpha))$ denotes the expected payoffs from the mixture α. As usual, this stage game has a Nash equilibrium. Let V be its set of feasible and individually rational payoffs.

We embellish the infinitely repeated game $G(\delta)$ with a correlation device that sends private messages to players each period conditional on the action history. The device makes public the private message profile *after* play concludes each period. Before each period (including the first), each player privately receives a message $\tilde{a}_i \in A_i$, which we interpret as a recommendation to play action a_i. By Aumann (1987), restricting messages to recommendations is without loss of generality.[3] Players commonly observe the null history $\mathbb{h}^1 = \emptyset$ before play begins. A history $\mathbb{h}^t = (a^1, \tilde{a}^1, \ldots, a^{t-1}, \tilde{a}^{t-1})$ is a complete record of all past *outcomes* in periods $1, 2, \ldots, t-1$, i.e. pairs of action and recommendation profiles. The history $\mathbb{h}^t$ is commonly observed by all players at the start of period t. Let $\mathbb{H}^t$ be the set of all histories $\mathbb{h}^t$, and $\mathbb{H} = \bigcup_{t=1}^{\infty} \mathbb{H}^t$ the set of all histories of any length.

A (direct) correlation device μ is a probability measure on the set of action profiles A. An *extensive form correlation device* is a sequence of functions $\lambda = (\lambda^t)_{t=1}^{\infty}$ such that $(\lambda^t : \mathbb{H}^t \to \Delta(A))_{t=1}^{\infty}$, and Λ is the space of all such functions.[4] The interpretation is that after history $\mathbb{h}^t$, the correlation device selects an action profile $\tilde{a} = (\tilde{a}_1, \cdots, \tilde{a}_n) \in A$ according to the distribution $\lambda(\mathbb{h}^t)$ and privately informs each player i of his recommended action $\tilde{a}_i$. Players then simultaneously choose actions. Finally, the recommendations are revealed to all players, and they become part

[2]Throughout, subscripts will denote players and superscripts will denote periods. Let $|X|$ denote the cardinality of X. Also, we parse any vector $x \equiv (x_i, x_{-i})$. Since we consider finite action and signal sets, all functions thereon are measurable.

[3]This can equivalently be justified by the Revelation Principle. In our finite model, the Revelation Principle holds since there cannot be issues with the measurable composition of functions.

[4]The notion of an extensive form correlation device is attributable to Forges (1986), who provided the canonical representation and geometric properties of extensive form correlation devices.

of the next history $\mathbb{h}^{t+1}$. Finally, let $G^\lambda(\delta)$ be the infinitely repeated mediated game with stage game G, extensive form correlation device $\lambda \in \Lambda$, and discount factor $0 < \delta < 1$.

A (behavior) strategy $\mathbb{s}_i$ for player i is a sequence $(\mathbb{s}_i^t)_{t=1}^\infty$, where $\mathbb{s}_i^t : \mathbb{H}_i^t \times A_i \to \Delta(A_i)$ for every period $t = 1, 2, \ldots$ So a strategy assigns a mixed action to every pair of history and recommendation. For any strategy profile $(s_1, \ldots, s_n) = \mathbb{s} \in \mathbb{S} = \prod_{i \in N} \mathbb{S}_i$, correlation device λ, and history $\mathbb{h}^t$, the payoff for player i is the present value of future payoffs:

$$\mathbb{v}_i^t(s|\mathbb{h}^t, \lambda) = (1-\delta)E\left[\sum_{r=t}^{\infty} \delta^{r-t} u_i(a^r) \middle| \lambda, \mathbb{s}, \mathbb{h}^t\right]$$

A strategy profile $\mathbb{s}$ is a sequential equilibrium of $G^\lambda(\delta)$ if in every period t, history $\mathbb{h}^t$, and alternative strategy $\tilde{\mathbb{s}}_i$,

$$\mathbb{v}_i^t(\mathbb{s}|\mathbb{h}^t, \lambda) \geq \mathbb{v}_i^t(\tilde{\mathbb{s}}_i, \mathbb{s}_{-i}|\mathbb{h}^t, \lambda)$$

1.4 Markov Perfect Correlated Equilibrium

If $\mathbb{s} \in \mathbb{S}$ is a sequential equilibrium strategy profile of $G^\lambda(\delta)$, then Prokopovych (2006) calls the pair $(\mathbb{s}, \lambda)$ a *Perfect Correlated Equilibrium (PCE)* of $G(\delta)$. The correlation device assumed in a PCE may depend arbitrarily on history. We now introduce a simpler solution concept that yields the same payoff prediction. A correlation device λ is *Markovian* if its recommendations depend solely on the outcome $(a, \tilde{a})$ of the most recent period. Denote by Λ_M the space of all such devices $\lambda : A^2 \to \Delta(A)$. Similarly, a strategy $\mathbb{s}$ is Markovian if it depends only on the most recent outcome and currently recommended action $\tilde{a}_i$, i.e. $\mathbb{s}_i : A^2 \times A_i$. If the device λ is Markovian, then there is a Markovian best response to a Markovian strategy (cf. Hernandez-Lerma, 1989 Theorem 2.2). Thus, a pair $(\mathbb{s}, \lambda)$ is a *Markov Perfect Correlated Equilibrium (MPCE)* of $G(\delta)$ if it is a PCE of $G(\delta)$ and both the correlation device λ and the strategy profile $\mathbb{s}$ are Markovian.

Let V^λ be the set of all sequential equilibrium payoff vectors of $G^\lambda(\delta)$. The *MPCE payoff set* V^* is the set of all payoff vectors attainable in an MPCE. Namely,

$$V^* \equiv \bigcup_{\lambda \in \Lambda_M} V^\lambda$$

The Appendix exploits self-generation methods to prove:

Lemma I.1. *Any PCE payoff is attainable in an MPCE.*

Because every MPCE is a PCE by definition, Lemma I.1 implies that both concepts yield the same equilibrium payoff sets.

Let $\mu \in \Delta(A)$ be a probability distribution on the set of action profiles A — as realized in a PCE as $\mu = \lambda(\mathbb{h}^t)$, or in an MPCE as $\mu = \lambda(a, \tilde{a})$. Fix a compact convex set of payoff vectors $W \subset \mathbb{R}^n$. A *continuation value function* $\mathbb{k} : A^2 \to W$ describes discounted future (equilibrium) payoffs for each current period outcome. Given the stage game payoffs, the mapping $\mathbb{k}$ completely describes the *auxiliary game* $G_{\mathbb{k}}$. This game is (the agent normal form of) a one-shot Bayesian game whose type profile $(\tilde{a}_1, \ldots, \tilde{a}_n) \in A$ is drawn from the distribution μ. Each player's type $\tilde{a}_i$ has the action set A_i, but the revised payoff function $E_\mu \left[(1-\delta)u_i(a) + \delta \mathbb{k}_i(a, \tilde{a}) | \tilde{a}_i\right]$ for the recommended action $\tilde{a}_i$.

If the distribution μ is a correlated equilibrium of $G_{\mathbb{k}}$, then the pair $(\mu, \mathbb{k})$ is *admissible w.r.t.* W, where W is the co-domain of $\mathbb{k}$. In this case,

$$E_\mu[(1-\delta)u_i(a) + \delta \mathbb{k}_i(a, \tilde{a}) | \tilde{a}_i] \geq E_\mu[(1-\delta)u_i(a'_i, a_{-i}) + \delta \mathbb{k}_i(a'_i, a_{-i}, \tilde{a}) | \tilde{a}_i] \tag{1.1}$$

for all players i, actions $a'_i \in A_i$, and recommendations $\tilde{a}_i \in A_i$ and $\tilde{a} \in A$. The *value* w of a pair $(\mu, \mathbb{k})$ is the (ex-ante) expected payoff $E_\mu[(1-\delta)u(a) + \delta \mathbb{k}(a, a)]$. Inversely, we write that the admissible pair $(\mu, \mathbb{k})$ *enforces* the payoff w on the set W if w is the value of the pair, and W is the co-domain of $\mathbb{k}$.

Let the set $B(W)$ be the union of all payoffs enforced on W, so that

$$B(W) = \left\{v = E_\mu[(1-\delta)\pi(a) + \delta \mathbb{k}(a, \tilde{a})] \,\middle|\, (\mu, \mathbb{k}) \text{ is admissible w.r.t. } W\right\}$$

Equivalently, $B(W)$ is the union of all correlated equilibrium payoffs in the auxiliary game $G_{\mathbb{k}}$, as $\mathbb{k}$ ranges over all continuation value functions with co-domain W.

The operator $B(\cdot)$ has some convenient properties. First, it is monotone: If $W \subseteq W'$, then $B(W) \subseteq B(W')$. Intuitively, the right side consists of the correlated equilibria of a larger set of auxiliary games. Secondly, $B(\cdot)$ is convex-valued: If $(\mu^1, \mathbb{k}^1)$ supports w^1 and $(\mu^2, \mathbb{k}^2)$ supports w^2, then for all weights $\theta \in [0, 1]$, the payoff $\theta w^1 + (1-\theta)w^2$ is supported by $(\theta\mu^1 + (1-\theta)\mu^2, \theta\mathbb{k}^1 + (1-\theta)\mathbb{k}^2)$.

As usual, we call a set $W \subset \mathbb{R}^n$ is *self-generating* if $W \subseteq B(W)$.

Theorem I.1 (MPCE Payoffs)**.** *The MPCE payoff set V^* has the properties:*

(a) It is the largest fixed point of $B(\cdot)$.

(b) It is a compact convex subset of V.

(c) It contains the convex hull of the set of SPE payoffs of $G(\delta)$

(d) It is nondecreasing in δ.

The proof is in the Appendix, but here we offer some intuition. First, part (a) captures the recursive structure of MPCE, which is analogous to factorization of PPE. If a set W is self-generating, then there exists an admissible pair with co-domain W. For any $w \in W$, a sequential equilibrium with payoff w can be constructed period-by-period by replacing every continuation value with a pair admissible w.r.t. W. This is always possible since W is self-generating.

Next, compactness in (b) follows since weak inequalities define incentive compatibility. Public randomization can always be created using a correlation device, and so the MPCE payoff set is convex. To publicly randomize between outcomes, let us step outside the space of direct devices and consider a new device that generates two messages for each player: the original message and a second that indicates the outcome of the public randomization. By the Revelation Principle, there exists an equivalent direct device.

For insight into part (c), consider the extensive form correlation device that recommends the subgame perfect equilibrium behavior after every history. By construction, this device constitutes a PCE, and Lemma I.1 guarantees that this payoff is attainable in an MPCE. Part (c) in particular implies that the folk theorem holds for MPCE.

Part (d) follows from the well-known principle that dynamic incentives can induce any behavior in patient players that it can in their less patient counterparts.

The MPCE payoff set can be obtained by iterating the B operator on a seed set $W^0 \subseteq \mathbb{R}^n$ containing the feasible and individually rational payoffs V. The algorithm starts by observing that $V^* \subseteq V \subseteq W^0$. Then either W^0 is self-generating or $B(W^0) \subseteq W^0$. Repeatedly applying $B(\cdot)$ to the inequality $V^* \subseteq W^k$, where $W^k = B(W^{k-1})$, produces a strictly decreasing sequence of nested sets that converges to the MPCE set V^*.

Theorem I.2 (Algorithm)**.** *The MPCE payoff set is $V^* = \lim_{j \to \infty} W^j$, where the payoff set W^0 obeys $V^* \subseteq W^0$, and $W^{j+1} = B(W^j)$ for $j = 1, 2, 3, \ldots$*

To implement the algorithm, we employ methods similar to those introduced by Judd, Yeltekin, and Conklin (2003). Compactness and convexity allow us to represent a set by its extreme points, and they imply that $B(W) = B(\text{ext } W)$. This makes the algorithm computationally tractable. The white area is the SPE payoff set; MPCE payoffs also include the grey area, so that these are MPCE payoffs unattainable in an SPE; the black area represents feasible and individually rational payoffs that are not MPCE, and thus unattainable in any sequential equilibrium.

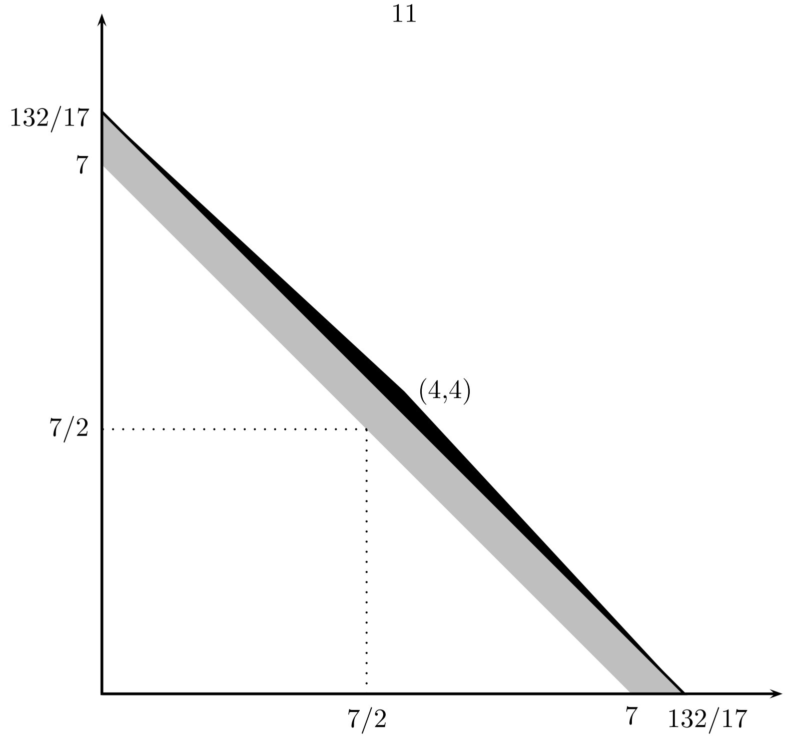

Figure 1.2: Payoffs in the Repeated Game in Figure 1.1.

Let's return to the repeated game of Section 1.2. In Figure 1.2, one can see that the MPCE payoff set is significantly larger than that of subgame perfect equilibrium. The extreme feasible and individually rational payoffs $(132/17, 0)$ and $(0, 132/17)$ are also the highest single player payoff vectors. So by convexity, the symmetric payoff $(66/17, 66/17)$ is also an MPCE, and in fact the highest symmetric MPCE payoff. This payoff is a convex combination of two extremal MPCE payoffs.

We now justify these claims. First, let us construct the device that delivers the highest payoff to one player. Let $(p, q, r, 1-p-q-r) \in \Delta(A)$ be the chances of $\{(C, D), (C, D), (C, D), (C, D)\}$, respectively, and $w_1, w_2 \in \mathbb{R}^2$ the continuation payoffs for players $1, 2$. Given the stage game of Figure 1.1, the highest MPCE payoff for player 1 solves

$$\max_{p,q,r,(w_1,w_2)\in V} (1-\delta)(4p - 13q + 20r) + \delta w_1$$

given: (i) $p, q, r \geq 0$ and $p + q + r \leq 1$, and (ii) payoffs are feasible and individually rational, and in particular $0 \leq w_1, w_2 \leq 132/17$, and (iii) two self-generation feasibility constraints that players not be promised payoffs higher than can be delivered:

$$w_1 \leq (1-\delta)(4p - 13q + 20r) + \delta w_1 \quad \text{and} \quad w_2 \leq (1-\delta)(4p + 20q - 13r) + \delta w_2$$

and (iv) two incentive constraints, for when players are told to play C:

$$(1-\delta)(4p-13q)+\delta w_1 \geq (1-\delta)20p \quad \text{and} \quad (1-\delta)(4p-13r)+\delta w_2 \geq (1-\delta)20r$$

Solving this program yields

$$132/17 = w_1 = 4p - 13q + 20r \quad \text{and} \quad 0 = 4p + 20q - 13r \quad \text{and} \quad p+q+r=1$$

So $(p,q,r) = (13/17, 0, 4/17)$. Then the payoff $(132/17, 0)$ is attainable in an MPCE. By symmetry, so too is the payoff $(0, 132/17)$. By convexity, the payoff $(66/17, 66/17)$ is an MPCE.

One can verify that imposing symmetry of the form $q = r$ yields a lower constrained maximum — i.e. a symmetric device does not yield the highest symmetric payoff. This implies that $(66/17, 66/17)$ is the highest symmetric MPCE payoff.

This effect is not limited to this example. In fact, a sufficient condition for correlation to be helpful in an infinitely repeated prisoner's dilemma is that: (i) mutual cooperation is efficient but not a subgame perfect equilibrium outcome, and (ii) the gain to defecting is higher when the other player cooperates than when he defects.

1.5 Repeated Games of Private Monitoring

A. The Stage Game. The structure here is standard, following closely the set-up of Ely, Horner, and Olszewski (2005). As in Section 1.3, a repeated game is played in periods $1, 2, \ldots$ Each period, every player $i \in N = \{1, 2, \ldots, n\}$ chooses an action a_i from a finite action set A_i. But now, after play any period, each player receives a private message m_i from a finite set M_i. A *monitoring structure* ψ is a collection of $|A|$ probability distributions $\{\psi(\cdot|a) \in \Delta(M) \,|\, a \in A\}$ on the message profile set $M = \prod_i M_i$. Let the set of all monitoring structures be Ψ. After an action profile a is realized, a message profile $m = (m_1, \ldots, m_n)$ is drawn with chance $\psi(m|a)$, and each player i is then privately informed of his component message m_i.

A player's realized payoff $\pi_i(a_i, m_i)$ following action a_i and message m_i depends on the other actions only through their effect on the private messages. In other words, observing one's payoff does not confer additional information. Player i's expected payoff from the action profile a is then

$$u_i(a) = \sum_{m_i \in M_i} \psi_i(m_i|a)\pi_i(a_i, m_i) \tag{1.2}$$

We shall consider different monitoring structures ψ consistent with the same "expected stage game". This requires that the payoffs $u(a) = (u_1(a), \ldots, u_n(a))$ not depend on the monitoring structure. Since payoffs depend on ψ in (1.2), this exercise implies a corresponding change in the stochastic payoff structure π. Such a choice is possible provided (1.2) is solvable in π_i for any ψ_i, and for all players i. This is feasible if and only if the matrix $(\psi_i(m_i|a_i, a_{-i}), m_i \in M_i, a_{-i} \in A_{-i})$ has full rank for every player i, and every action a_i. This requires that each player can statistically identify the actions of his opponents.[5] This generically holds when, for instance, everyone has at least as many messages as there are players. We assume that this condition is met by any monitoring structure in Ψ under consideration. Our results do not explicitly depend on this; it simply allows us to meaningfully consider a fixed stage game.

B. The Repeated Game. Let $G_\psi(\delta)$ denote the infinitely repeated game of private monitoring with monitoring structure ψ, played in periods $t = 1, 2, 3, \ldots$. Payoffs are discounted as usual by the factor $0 < \delta < 1$. The game reduces to a standard repeated game with perfect monitoring when private messages are action profiles, i.e. if $M_i = A$ and $\psi_i(m_i|a) = 1$ when $m_i = a$ and 0 otherwise, for all players i. Similarly, the game reduces to a standard repeated game with public monitoring if $M_i = M$ for all players i, and $\psi_i(m|a) = 1$ if and only if $\psi_j(m|a) = 1$ for every pair of players i, j.

In each period, a player observes his realized action $a_i \in A_i$ and private message m_i. Let the null history h_i^1 be player i's history before play begins. A *private history* h_i^t is the complete record of player i's past actions $(a_i^1, \ldots, a_i^{t-1})$ and past private messages $(m_i^1, \ldots, m_i^{t-1})$, including the null history. Let H_i^t be the set of all possible private histories h_i^t for player i, and $H_i = \bigcup_{t=1}^{\infty} H_i^t$ the set of all such histories of any length. A (behavior) strategy s_i is a sequence of functions $\{s_i^t\}_{t=1}^{\infty}$, where $s_i^t : H_i^t \to \Delta(A_i)$ for every period $t = 1, 2, 3, \ldots$ In other words, it maps every private into a mixed action. Let $\mathcal{S}$ be the space of all such strategy profiles $s = (s_1, \ldots, s_n)$.

Given the strategy profile $s \in \mathcal{S}$, Bayes' rule and the Law of Total Probability naturally imply beliefs and behavior at all future information sets. Let $v_i : \mathcal{S} \to \mathbb{R}$ be the discounted average payoff for player i in the repeated game $G_\psi(\delta)$. While more precisely presented in the Appendix, here we write that player i's discounted average payoff starting in period t from the strategy profile s is $v_i^t(s|h_i^t)$. Then a strategy profile s is a *sequential equilibrium* of $G_\psi(\delta)$ if and only if no player can ever profitably deviate, i.e. $v_i(s|h_i^t) \geq v_i(\tilde{s}_i, s_{-i}|h_i^t)$ for every private history h_i^t and strategy $\tilde{s}_i : H_i \to \Delta(A_i)$ of every player i. Since playing a Nash equilibrium of G after every history is a sequential equilibrium, existence is guaranteed. Let V_ψ be the set of sequential

[5] This is somewhat analogous to the pairwise full rank condition of Fudenberg, Levine, and Maskin (1994), which requires that each player be able to statistically identify the actions of another player.

equilibrium payoff vectors of the mediated game $G_\psi(\delta)$.

1.6 Unattainable Private Monitoring Payoffs

A. An Upper Bound. We bound the sequential equilibrium payoffs by the MPCE payoff set V^*. This inclusion might at first blush appear surprising: For the repeated game $G_\psi(\delta)$ has no proper subgames, whereas $G^\lambda(\delta)$ introduces a new subgame every period. So while continuation play in $G^\lambda(\delta)$ is common knowledge, it is not so in $G_\psi(\delta)$. We proceed by associating outcomes in $G_\psi(\delta)$ with those of $G^\lambda(\delta)$. To do so, we replace the endogenous correlated beliefs in $G_\psi(\delta)$ with those from a fixed correlation device λ. Also, we do so in an incentive compatible fashion.

Theorem I.3 (Upper Bound)**.** *For any monitoring structure ψ, every sequential equilibrium payoff of the repeated game $G_\psi(\delta)$ is attained in an MPCE of $G(\delta)$.*

This implies that MPCE captures the payoffs in many studied subclasses of equilibria. It contains all PPE payoffs for any public monitoring structure, as well as all sequential equilibrium payoffs in private strategies (Kandori and Obara, 2006), as well as all belief-free and weakly-belief-free equilibrium payoffs (Kandori, 2008).

The proof in the Appendix first deduces this result for PCE, and then appeals to Lemma I.1. The proof for PCE involves two steps. We show that for any strategy profile $s \in \mathcal{S}$, there exists a correlation device $\lambda \in \Lambda$ and strategy profile $\mathbb{s} \in \mathbb{S}$ that induce in $G^\lambda(\delta)$ the same outcome as does s in $G_\psi(\delta)$. After the history $\mathbb{h}^t$ in the mediated game $G^\lambda(\delta)$, the correlation device draws a "fictitious private history" h_i^t for each player $i \in N$ according to the true posterior probability of that history conditional on the actions of history $\mathbb{h}^t$. The device then recommends the actions prescribed at that private history profile h^t by the continuation strategy profile $s(h^t)$. By induction on the period t, we show that the distribution over recommendations in the mediated game coincides with the distribution of actions in $G_\psi(\delta)$. In our next step, we argue that if s is a sequential equilibrium strategy profile of $G_\psi(\delta)$, then λ constitutes a PCE. For if some player has a profitable deviation in $G^\lambda(\delta)$, then we argue that he must also have one in $G_\psi(\delta)$. The argument turns on the equivalence of beliefs about continuation play in $G^\lambda(\delta)$ and $G_\psi(\delta)$.

B. A Tight Upper Bound. Since this upper bound is independent of the monitoring structure ψ, one might think that the inclusion in Theorem I.3 could not be tight. In fact, this is true, but only because correlated play in a private monitoring game starts no earlier than the second period. So inspired, we now exploit the MPCE payoffs to deduce a tight upper bound for equilibrium payoffs of private monitoring games.

For a standard repeated game played in periods $1, 2, 3, \ldots$, we can remove the first period correlation from MPCE. An admissible pair $(\mu, \Bbbk)$ is called *Nash admissible* if μ is the result of independent mixtures, i.e. $\mu \in \prod_i \Delta(A_i)$. We then obtain the operator from APS, here denoted by B_{NE}:

$$B_{NE}(W) = \left\{ v = E_\mu[(1-\delta)\pi(a) + \delta \Bbbk(a, \tilde{a})] \,\middle|\, (\mu, \Bbbk) \text{ is Nash admissible w.r.t. } W \right\}$$

This collects the Nash equilibrium payoffs of all auxiliary games formed with continuation value functions mapping into W. Since first period strategies are uncorrelated in $G_\psi(\delta)$, we use a two-stage procedure. First, we compute the MPCE payoff set, and then use this set $W = V^*$ as continuation payoffs in $B_{NE}(W)$.

Theorem I.4 (Tightness)**.** *A payoff is Nash admissible w.r.t. the MPCE set of $G(\delta)$ if and only if it is a sequential equilibrium payoff of $G_\psi(\delta)$ for some monitoring structure ψ, so that*

$$\bigcup_{\psi \in \Psi} V_\psi = B_{NE}(V^*)$$

Without reference to the monitoring structure, there exists no tighter bound on the sequential equilibrium payoffs in a repeated game of private monitoring.

In the example of Section 1.2, Theorem I.3 demonstrates that $(66/17, 66/17)$ is the highest symmetric sequential equilibrium payoff in the infinitely repeated game *with any monitoring structure*, and so all symmetric payoffs in $(66/17, 4]$ are unattainable. In fact, except for the payoffs $(132/17, 0)$ and $(0, 132/17)$, all efficient payoff vectors are unattainable in a sequential equilibrium.

C. How Restrictive Is Perfect Public Equilibrium? We are now in a place to discuss what is lost by restricting attention to public monitoring. Since the set of PPE payoffs is a subset of subgame perfect payoffs, the demonstrated gap between the MPCE and SPE payoff sets implies that a public solution concept may fail to capture the potential outcomes of environments with richer information in which the folk theorem is silent. Thus a regulator attempting to detect antitrust violations may, upon observing payoffs inconsistent with some PPE, draw the wrong conclusion about collusion. These efforts ought to keep in mind the strategic opportunities afforded by complex information structures; This is done precisely by using MPCE in the place of a public monitoring solution concept. Furthermore, in many applications the relevant monitoring structure is difficult to determine, and thus PPE is difficult to use. Thus unlike PPE, MPCE enables one to study equilibrium payoffs while being agnostic about the monitoring structure.

We will now demonstrate that for a generic class of prisoner's dilemma games, if the discount factor is high, but not too high, correlation improves upon subgame perfect equilibrium, and hence perfect public equilibrium. Consider the infinite repetition of the following prisoner's dilemma, where players share the discount factor δ.

	C	D
C	(1,1)	(-c,b)
D	(b,-c)	(0,0)

Figure 1.3: Prisoner's Dilemma ($b > 1, c > 0$)

As before, we assume that (C, C) is the efficient action profile (i.e. $b - c < 2$), and that the gain to defecting when the opponent plays C is larger than the gain when he plays D (i.e. $b - 1 < c$). Stahl (1991) characterizes the subgame perfect equilibrium payoff correspondence with respect to the discount factor. If $\delta < \frac{c}{b}$ then the only repeated game payoff vector is the stage game Nash equilibrium payoff $(0, 0)$. When $\delta \geq \frac{b-1}{b}$, every feasible and individually rational payoff vector is a subgame perfect equilibrium payoff vector. However, if $\delta \in [\frac{c}{b}, \frac{b-1}{b})$ then the set of subgame perfect payoffs is the triangle $T = \{(0, 0), (b - c, 0), (0, b - c)\}$.

When the discount factor is in this intermediate range, correlation can be used to support a payoff vector $(v^*, v^*) \notin T$. To see this, consider the continuation value function

$$\Bbbk(a, \tilde{a}) = \begin{cases} (v^*, v^*) & \text{if } a = \tilde{a} = (C, C) \\ (b - c, 0) & \text{if } a = \tilde{a} = (D, C) \\ (0, b - c) & \text{if } a = \tilde{a} = (C, D) \\ (0, 0) & \text{if } a = \tilde{a} = (D, D) \\ (0, 0) & \text{if } a \neq \tilde{a} \end{cases}$$

We can then check whether the auxiliary game implied by this continuation value function takes the form of "chicken", and therefore has non-trivial correlated equilibria. For the game to be "chicken" it must be that

$$(1 - \delta) + \delta v^* \leq (1 - \delta)b$$

$$(1 - \delta)(-c) + \delta(b - c) \geq 0$$

Both expressions are satisfied for a payoff vector $v^* \notin T$ over the entire range of parameters considered.

1.7 Conclusion

Understanding the equilibria of repeated games with private monitoring has long been the next frontier in game theory. Yet finding sequential equilibria here has been hard, because existing recursive methods only capture subsets of them. In this paper, we have developed a new solution concept for repeated games, Markov Perfect Correlated Equilibrium, with a recursively computable payoff set. This is the smallest set that contains all equilibrium payoffs of the analogous repeated game endowed with any monitoring structure. It therefore provides insights into important economic environments while being agnostic about specific, possibly unobservable, informational aspects of the game. We also hope our bound will offer a rebirth to the recursive methods of Abreu, Pearce, and Stacchetti (1990) in settings with richer information structures than they had envisioned.

CHAPTER II

Strategically Valuable Information

2.1 Introduction

In his classic 1953 paper, Blackwell compares two partial orderings on experiments, or informative signals. The first ranks experiments by statistical sufficiency — experiment A is sufficient for B if B is a statistical garbling of A, or intuitively, B can be attained by adding noise to A. The second ranks experiments by their economic value in decision problems with state-dependent payoffs. Experiment A is more valuable than B if the decision maker can attain a larger set of payoffs with A than with B in any decision problem. Blackwell showed that experiment A is sufficient for B if and only if it is more valuable.

In this paper, we ask the Blackwell equivalence question in a strategic setting. In other words, does there exist a partial order on information held by players in a game that reflects "more" or "better" information, which coincides precisely with the ability to induce more equilibrium payoff vectors in all Bayesian games? If so, we say that it is "strategically more valuable". In this paper, we define a meaningful sense in which information structures can be compared by how "strategically informative" they are. Combining the two notions, we answer our original question in the affirmative: There exists an intuitive definition and characterization of the partial order *more strategically informative*, and it is equivalent to the partial order *more strategically valuable.* The conditions we provide are easily checked, are useful in an array of economic settings, and have straightforward geometric interpretations.

Our main theorem applies to a wide variety of economic environments of interest endowed with commonly used information structures. For example, sunspots are a frequently used tool in general equilibrium theory. Our results provide a natural partial ordering on sunspot equilibria, regardless of the environment in which they operate. The centerpiece application is to repeated games with private monitoring. In a repeated game with private monitoring, each period players simultaneously choose actions, after which each player privately observes a signal informative of the action profile most recently played. Neither the actions nor signal realizations are ever observed by any other player. The more strategically informative order ranks monitoring structures — the probability distribution on private signals — in much the same way. Consequently, we can show when a change in monitoring structure will weakly expand the set of sequential equilibria. This mirrors a classic result of Kandori (1992) for repeated games with imperfect public monitoring.

The most closely related paper to this one is that by Gossner (2000). He studies the same question, and succeeds in obtaining a different characterization. We feel that the current approach is superior in several dimensions. First, our condition is easier to verify — it corresponds geometrically to well understood statistical concepts. Second, while the proofs in Gossner (2000) are indirect and complex, we provide a straightforward, illuminating proof that leverages the separation argument at the core of Blackwell's Theorem.

2.2 Model

2.2.1 The Informational Setting

We begin with the "extrinsic uncertainty" of our environment, the probability space $\mathbb{P} = (\Omega, \mathcal{F}, \mu)$. The set Ω represents the sample space, i.e. all possible states of the world. Measurable events are described by $\mathcal{F}$, a σ-algebra of Ω. Finally, μ is a (prior) measure on the events of $\mathcal{F}$.

The environment is populated with $n \geq 2$ players. Each player's private information is described by a (possibly infinite) *partition* $\mathcal{P}_i$ of Ω, a mutually exclusive and exhaustive family of subsets of Ω measurable with respect to $\mathcal{F}$. At a state $\omega \in \Omega$ the set $p_i(\omega) \in \mathcal{P}_i$ is the set of states indistinguishable by player i, which we will call a *cell.* Similarly, let $\mathcal{P} = \{\mathcal{P}_i\}_{i=1}^n$ be the *joint partition* of all players.

An *information structure* $\mathcal{I}$ is a pair $(\mathbb{P}, \{\mathcal{P}_i\}_{i=1}^n)$ of a probability space $\mathbb{P}$ and a joint partition $\mathcal{P} = \{\mathcal{P}_i\}_{i=1}^n$. Assuming that the relevant probability space is $\Omega = [0, 1]$ endowed with the Borel σ-algebra and the Lebesgue measure entails no loss of generality, so in much of what follows we do so. In that case, an information structure is completely described by the joint partition $\mathcal{P}$. Consider the following two player information structures:

$$\mathcal{P}_1 = \{[0, 1/3), [1/3, 1]\} \quad \mathcal{P}_2 = \{[0, 2/3), [2/3, 1]\}$$

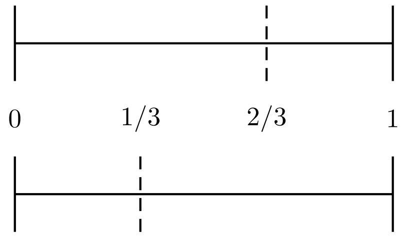

Figure 2.1: Information Structure $\mathcal{P}$

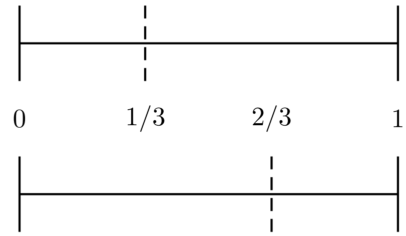

Figure 2.2: Information Structure $\mathcal{Q}$

and

$$\mathcal{Q}_1 = \{[0, 2/3), [2/3, 1]\} \quad \mathcal{Q}_2 = \{[0, 1/3), [1/3, 1]\}$$

The information structure $\mathcal{Q}$ is obtained from $\mathcal{P}$ by 'relabeling', in this case reversing the direction of the interval $[0, 1]$. Since the state is not payoff relevant, it is natural to require that these information structures have equivalent strategic effects in any environment. In fact, by the same logic any (measure preserving) permutation of the state space should be similarly inconsequential and for the purposes of this paper be treated as equivalent. Let $[\mathcal{P}]$ be the set of information structures equivalent to $\mathcal{P}$: $\mathcal{Q} \in [\mathcal{P}]$ if there exists a transformation t in the permutation group $T(\Omega)$ such that $\mathcal{P} = t(\mathcal{Q})$.[1] We can then define the equivalence relation $\sim$: $\mathcal{P} \sim \mathcal{Q}$ if and only if $\mathcal{Q} \in [\mathcal{P}]$. This relation is reflexive, symmetric and transitive. In what follows, we will work on the quotient space of information structures modulo $\sim$, defined by the canonical projection from the space of all information structures to the quotient space. Therefore, without loss of generality, we will not distinguish between elements of an equivalence class.

Let $\sigma(\mathcal{P}_i)$ be the σ-algebra generated by $\mathcal{P}_i$, and $\sigma(\mathcal{P})$ the σ-algebra generated by $\mathcal{P}$ on the product space $\times_{i=1}^{n}\Omega$. Information structure $\mathcal{P}$ *refines* $\mathcal{Q}$ if $\sigma(\mathcal{Q}_i) \subseteq \sigma(\mathcal{P}_i)$ for every player i. Two sub-σ-algebras $\mathcal{F}, \mathcal{G}$ are *conditionally independent given σ-algebra $\mathcal{H}$* — written $(\mathcal{F} \perp \mathcal{G})|\mathcal{H}$ — if any two events $F \in \mathcal{F}$ and $G \in \mathcal{G}$ are conditionally independent given $\mathcal{H}$, namely:

$$\mu(F \cap G|\mathcal{H}) = \mu(F|\mathcal{H})\mu(G|\mathcal{H})$$

Information structure $\mathcal{P}$ is *more strategically informative* than $\mathcal{Q}$ — written $\mathcal{P} \gg \mathcal{Q}$ — if $\mathcal{P}$ refines $\mathcal{Q}$ and $(\mathcal{P}_i \perp \mathcal{Q}_{-i})|\mathcal{Q}_i$ for every player i . This requires that the information added by $\mathcal{P}$ (relative to $\mathcal{Q}$) be safely ignored by any player: if players $-i$ ignore the added information, nothing is lost by player i by doing so as well. Conditional independence means that player i's

[1] A *permutation group* (of transformations) T is a set of measure preserving bijective automorphisms on a set X such that: (a) If $t_1 \in T$ and $t_2 \in T$, then $t_1 \circ t_2 \in T$; (b) For all transformations $t \in T$ there exists $t^{-1} \in T$ such that for all elements $x \in X$we have $(t^{-1} \circ t)(x) = x$, i.e. there exists an inverse; and (c) There exists $e \in T$ such that $e(x) = x$, i.e. there exists an identity.

information $\mathcal{Q}_i$ is sufficient for $\mathcal{Q}_{-i}$. In other words it does not help player i make inferences about the other players.

Example 1

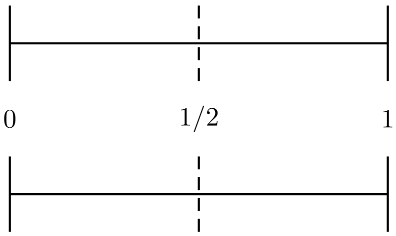

Figure 2.3: Sunpot with Two Ouctomes

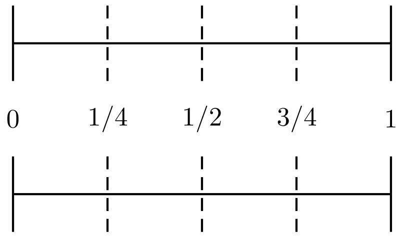

Figure 2.4: Sunspot with Four Outcomes

Consider the two-player information structures depicted in Figures 2.3 and 2.4:

$$\mathcal{P}_1 = \mathcal{P}_2 = \{[0, 1/2), [1/2, 1]\}$$

and

$$\mathcal{Q}_1 = \mathcal{Q}_2 = \{[0, 1/4), [1/4, 1/2), [1/2, 3/4), [3/4, 1]\}$$

In this example, $\mathcal{P}$ is a sunspot with two outcomes and $\mathcal{Q}$ is a sunspot with four outcomes. $\mathcal{P}$ is strategically more informative than $\mathcal{Q}$. □

Example 2

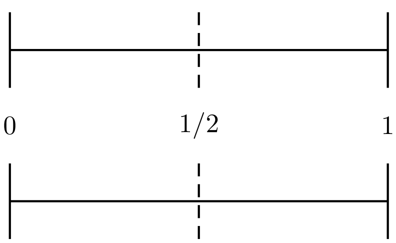

Figure 2.5: Information Structure $\mathcal{P}$

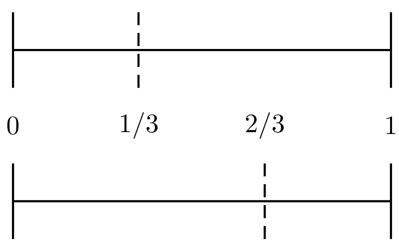

Figure 2.6: Information Structure $\mathcal{Q}$

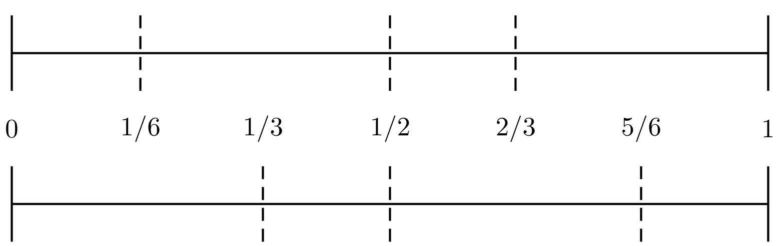

Figure 2.7: Information Structure $\mathcal{R}$ is constructed by embedding $\mathcal{Q}$ in both halves of $\mathcal{P}$

Now, consider the following three information structures, each for two players:

$$\mathcal{P}_1 = \mathcal{P}_2 = \{[0, 1/2), [1/2, 1]\}$$

$$\mathcal{Q}_1 = \{[0, 1/3), [1/3, 1]\} \quad \mathcal{Q}_2 = \{[0, 2/3), [2/3, 1]\}$$

$$\mathcal{R}_1 = \{[0, 1/6), [1/6, 1/2), [1/2, 2/3), [2/3, 1]\} \quad \mathcal{R}_2 = \{[0, 1/3), [1/3, 1/2), [1/2, 5/6), [5/6, 1]\}$$

In this example, $\mathcal{P}$ is again a sunspot with two outcomes. The information structure $\mathcal{Q}$ is the signal used in Aumann's original paper on correlated equilibrium Aumann (1974). Neither $\mathcal{P} \gg \mathcal{Q}$ nor $\mathcal{Q} \gg \mathcal{P}$. However, $\mathcal{R} \gg \mathcal{P}$ and $\mathcal{R} \gg \mathcal{Q}$.

2.2.2 Formulation as Signals

Example 3: Bivariate Gaussian Signals

Suppose there are two players who each observe a signal generated by a bivariate normal distribution with variance normalized to 1 and covariance ρ. With a slight abuse of notation, let $\mathcal{P}_\rho$ refer to the information structure induced by bivariate Gaussian signals with covariance ρ. Any correlation is more strategically informative than independent signals, that is if $\rho \neq 0$ we have $\mathcal{Q}_\rho \gg \mathcal{Q}_{\mathbf{0}}$, where $\mathbf{0}$ denotes independence. It should also be clear that by a simple transformation $\mathcal{Q}_\rho$ and $\mathcal{Q}_{-\rho}$ are equally strategically informative. However, for distinct non-zero covariances p, q inducing information structures $\mathcal{P}$ and $\mathcal{Q}$, respectively, neither $\mathcal{P} \gg \mathcal{Q}$ nor $\mathcal{Q} \gg \mathcal{P}$.

Example 4: Multivariate Gaussian Signals

When the information structure is composed of normally distributed signals, even when there are n players, the more strategically informative ordering has a particularly convenient representation. Suppose there is a vector valued random variable $X \sim N(0, \Sigma_X)$, where Σ_X is an $n \times n$ positive semi-definite (covariance) matrix. Then the state space $\Omega = \mathbb{R}^n$. Each of n players is informed of the projection of X onto Ω_i, i.e. the i-th component X_i of X.

Suppose the random variable $Y \sim N(0, \Sigma_Y)$. Let Σ_{XY} be the covariance matrix for the random variable $[X, Y]$. Let M_i be an $(n+1) \times (n+1)$ matrix obtained by eliminating rows and columns 1 to $i-1$ and $i+1$ to n. M_i is the covariance matrix for player i's information under X and all players information under Y. By a known characterization of conditional independence for multivariate Gaussian distributions, we have the following result. *The random variable X is more strategically informative than Y if and only if M_i is singular for every player i.*

Application: Repeated sampling So far we have considered a setting in which each agent observes his component of one signal. In many interesting economic environments agents observe multiple signals. Therefore, we extend our results to repeated signals, showing they are more strategically valuable.

Suppose instead of observing a realization of one signal, each player observes $K > 1$ signals. Now, nature chooses a vector $(\omega^1, \ldots, \omega^K)$ corresponding to the realization of K independent random variables. For each random variable $k \leq K$ there is a probability space $\mathbb{P}^k = (\Omega^k, \mathcal{F}^k, \mu^k)$. A player's private information is described a (possibly infinite) partition $\mathcal{P}_i$ of $\prod_k \Omega^k$, a mutually exclusive and exhaustive family of subsets of $\prod_k \Omega^k$ measurable with respect to the product σ-algebra $\prod_k \mathcal{F}^k$. Since for $j \neq k$ we have $\mathcal{F}^j \perp \mathcal{F}^k$, it is clear that more signals increases strategic information.

2.2.3 The Strategic Environment

We now turn to the strategic aspects of these environments. An n-player game G in normal form is a set of players $N = \{1, 2, \ldots, n\}$, a set A_i of pure actions available to each player i, and a function g from the set of action profiles $A = \prod_i (A_i)$ into $\mathbb{R}^n$. Let Γ_n be the set of all n-player games in normal form. A game $G \in \Gamma_n$ *extended by* $\mathcal{P}$ allows each player to choose his action conditional on the cell of his partition in which he finds himself. Therefore, a strategy is now a function f_i from $\mathcal{P}_i$ to A_i, where $f_i(p_i)$ is the action played by player i in cell p_i. The expected payoff function $E[g(f(p)|\mathcal{P})]$ is evaluated in the obvious way by using Bayes rule.

A strategy profile f is a *Bayes-Nash equilibrium* if

$$E[g_i(f(p)|\mathcal{P})] \geq E[g_i(f_i'(p_i), f_{-i}(p_{-i})|\mathcal{P})]$$

for all players i and all strategies f_i'. If f is a Bayes-Nash equilibrium, then $E[g(f(p)|\mathcal{P})]$ is a Bayes-Nash equilibrium payoff. Let $\Pi(G|\mathcal{P})$ be all Bayes-Nash equilibrium payoffs in the game G extended by information structure $\mathcal{P}$.

An information structure $\mathcal{P}$ is *more strategically valuable* than information structure $\mathcal{Q}$ — written $\mathcal{P} \sqsupseteq \mathcal{Q}$ — if $\Pi(G|\mathcal{P}) \supseteq \Pi(G|\mathcal{Q})$ for all games $G \in \Gamma_n$.

2.3 The Equivalence Result

We are now in a position to state the main result of this paper, which we prove in the appendix.

Theorem II.1. *Information structure $\mathcal{P}$ is more strategically valuable than information structure $\mathcal{Q}$ if and only if $\mathcal{P}$ is more strategically informative than $\mathcal{Q}$:*

$$\mathcal{P} \sqsupseteq \mathcal{Q} \Leftrightarrow \mathcal{P} \gg \mathcal{Q}$$

We now return to our previous example to demonstrate the implications of Theorem II.1. As we showed, neither $\mathcal{P}$ nor $\mathcal{Q}$ is more strategically informative than the other. Theorem II.1 implies that they are also not ranked by the more strategically valuable relation. To see this, we will demonstrate a game in which information structure $\mathcal{P}$ induces an equilibrium payoff not attained with $\mathcal{Q}$, and vice versa. In the coordination game depicted in Figure 2.8, the signal provided by $\mathcal{P}$ generates the payoff $(3/2, 3/2)$. This payoff requires perfect coordination and so cannot be obtained with $\mathcal{Q}$. Similarly, the information structure $\mathcal{Q}$ induces the payoff $(10/3, 10/3)$ in the game of chicken, depicted in Figure 2.9. This payoff cannot be attained with the sunspot.

	L	R
U	(2,2)	(0,0)
D	(0,0)	(1,1)

Figure 2.8: Pure Coordination Game

	L	R
U	(4,4)	(1,5)
D	(5,1)	(0,0)

Figure 2.9: Game of Chicken

2.4 Economic Applications

2.4.1 Application to Repeated Games

Kandori (1992) shows that in a repeated game with imperfect public monitoring, the equilibrium payoff set expands in the accuracy of the public signal. Our result extends this line of thought to repeated games with private monitoring. In particular, we show that the set of sequential equilibrium payoffs grows when the monitoring structure becomes more strategically informative.

A repeated game is played in periods $1, 2, \ldots$ Each period, each player $i \in N = \{1, 2, \ldots, n\}$ chooses an action a_i from a finite action set A_i. Players do not observe the actions of others, instead observing private signals of the period's outcome. After play any period, each player receives a private message m_i from a finite set M_i. A *monitoring structure* ψ is a collection of $|A|$ probability distributions $\{\psi(\cdot|a) \in \Delta(M) \,|\, a \in A\}$ on the message profile set $M = \prod_i M_i$. Let the set of all monitoring structures be Ψ. After an action profile a is realized, a message profile $m = (m_1, \ldots, m_n)$ is drawn with chance $\psi(m|a)$, and each player i is then privately informed of his component message m_i.

A *monitoring space* is a collection of $|A|$ probability spaces $\{\mathbb{P}(a)\}_{a \in A}$, and a joint partition of the sample space for each action profile. The more common notion of a monitoring structure consisting of a set of messages for each player and a probability distribution over message profiles for each action profile is clearly a special case of a monitoring space. For consistency with the literature, let each player's message space be his partition of the sample space and let the set of message profiles be the meet of all players' partitions.

In each period, a player knows his realized action $a_i \in A_i$ and observes his private message m_i. Let the null history h_i^1 be player i's history before play begins. A *private history* h_i^t is the complete record of player i's past actions $(a_i^1, \ldots, a_i^{t-1})$ and past private messages $(m_i^1, \ldots, m_i^{t-1})$, including the null history. Let H_i^t be the set of all possible private histories h_i^t for player i, and $H_i = \bigcup_{t=1}^{\infty} H_i^t$ the set of all such histories of any length. A (behavior) strategy s_i is a sequence of functions $\{s_i^t\}_{t=1}^{\infty}$, where $s_i^t : H_i^t \to \Delta(A_i)$ for every period $t = 1, 2, 3, \ldots$ In other words, it maps every private history into a mixed action. Let $\mathcal{S}$ be the space of all such strategy profiles $s = (s_1, \ldots, s_n)$. Given the strategy profile $s \in \mathcal{S}$, Bayes' rule and the Law of Total Probability naturally imply beliefs and behavior at all future information sets.

Each length t private history, together with a strategy profile, implies an ex-ante distribution on the product space $A^t \times M^t$. Each player, being informed only of his own actions and signals, entertains a natural partition of $A^t \times M^t$. For each period t, this information structure is defined endogenously by the distribution of mixed actions and monitoring signals. Let $\Phi^t(\psi, s)$ be the (ex-ante expected) information structure in period t under monitoring structure ψ and strategy profile s. A monitoring structure ψ^1 is *more strategically informative* than ψ^2 if

$$\left(\psi_i^2(\cdot|a) \perp \psi_{-i}^1(\cdot|a)\right) \bigm| \psi_i^1(\cdot|a)$$

for every player i and every action profile a. In a static context, this precisely coincides with the previous definition.

Lemma II.1. *If a monitoring structure ψ^1 is more strategically informative than ψ^2, then $\Phi^t(\psi^1, s)$ is more strategically informative than $\Phi^t(\psi^2, s)$ for any strategy profile s and any period t.*

Let $G_\psi(\delta)$ denote the infinitely repeated game of private monitoring with monitoring structure ψ, played in periods $t = 1, 2, 3, \ldots$. Payoffs are discounted as usual by the factor $0 < \delta < 1$. Let

$v_i : \mathcal{S} \to \mathbb{R}$ be the discounted average payoff for player i in the repeated game $G_\psi(\delta)$. While more precisely presented in the Appendix, here we write that player i's discounted average payoff starting in period t from the strategy profile s is $v_i^t(s|h_i^t)$. Then a strategy profile s is a *sequential equilibrium* of $G_\psi(\delta)$ if and only if no player can ever profitably deviate, i.e. $v_i(s|h_i^t) \geq v_i(\tilde{s}_i, s_{-i}|h_i^t)$ for every private history h_i^t and strategy $\tilde{s}_i : H_i \to \Delta(A_i)$ of every player i. Since playing a Nash equilibrium of G after every history is a sequential equilibrium, existence is guaranteed. Let V_ψ be the set of sequential equilibrium payoff vectors of the repeated game $G_\psi(\delta)$

Suppose monitoring structure ψ^1 is more strategically valuable than monitoring structure ψ^2. The main theorem implies that the sequential equilibrium repeated game payoff set with ψ^1 contains that of ψ^2. An immediate implication of the lemma is the following theorem.

Theorem II.2. *If a monitoring structure ψ^1 is more strategically informative than ψ^2, then $V_{\psi^1} \supseteq V_{\psi^2}$*

2.4.2 Sunspots in General Equilibrium

This result has important implications for general equilibrium theory. Since Debreu (1952) showed that Walrasian settings can be interpreted as games, our result applies to markets as well. Public signals that are not payoff relevant is often used in these environments to add convexity to outcomes. These signals, popularly known as sunspots, make goods divisible and help with equilibrium selection. Consequently, sunspots play a central role in many general equilibrium models.

We reinterpret competitive markets as games and model sunspots precisely with information structures. By doing so, the main result of this paper allows us to say with certainty when "better" public information forces the set of sunspot equilibria to grow. For example, as in Example 1, adding nested outcomes necessarily makes the set of equilibria (weakly) expand. However, suppose all agents observe the realization of a Gaussian random variable before acting. In this case the variance of the random variable is of no consequence; all sunspots with any positive variance are equally strategically valuable since observations can be transformed by any non-zero scalar.

2.5 Conclusion

At first glance, the main result should extend to games of incomplete information by introducing the player "Nature", as in the tradition of Harsanyi. However, doing so requires the same conditional independence conditions between each player and Nature as between any two players.

For intuition, consider the following simple example. There are two players $\{1, 2\}$, two states of the world $\{l, h\}$ and two pure actions for each player $\{L, H\}$. The game is zero-sum, with player 1 as the maximizer. He earns a payoff of 1 when his action matches the state, and zero otherwise. More precise information about the state increases his payoff, which necessarily lowers his opponent's payoff.

CHAPTER III

Electing Efficiency

3.1 Introduction

When every individual's effort imposes negative externalities on his competitors, competition results in excessive aggregate effort. Forming output-sharing groups among the competitors results in free riding which tends to reduce the aggregate effort. If groups of the right size happen to form, inefficiency can (in theory) be eliminated. This paper investigates in a controlled laboratory setting whether agents, given an opportunity to choose the size of their output-sharing groups, will eliminate or at least reduce the inefficiencies due to negative externalities.

In competitions to exploit common properties, aggregate effort by competing fishermen exceeds the socially efficient level (Dasgupta and Heal, 1985) —the so-called "tragedy of the commons." Similarly, in innovation tournaments, aggregate investment by firms competing to develop the best innovation exceeds the level that would maximize their aggregate profit (Baye and Hoppe, 2003)— the so-called "problem of R&D duplication." Moreover, in sports contests, aggregate investment by competing teams to identify and cultivate the athletic talent of their players goes beyond the level that would maximize total earnings (Canes 1974, Dietl, Franck, and Lang 2008 —the so-called "problem of ruinous competition between clubs." Henceforth, we will refer to such situations as the "negative aggregate spillover problem."

Since the competition among players who impose negative externalities on each other results in lost profit opportunities, arrangements have evolved to dampen effort incentives.[1] A common arrangement is for competitors to form groups which *share* revenue or output. Typically, effort costs are not shared because effort is too costly to monitor. Thus, as the collective work of Nobel

[1]According to anthropologists, those hunter-gatherers who have survived may have done so because their traditional practice of sharing the fish and game they caught deterred them from exhausting their resource base (Kagi 2001; Sahlins 1972).

Laureate Elinor Ostrom has documented (among others, Ostrom (1990), Ostrom and Walker (1991), and Ostrom, Walker, and Gardner (1994)), extractors of a common property often form groups which share whatever is extracted by any member of the group[2]; to avoid wasteful duplication, researchers often form joint ventures and share the benefits of their discoveries within their group (Kamien, Muller, and Zang, 1992); finally, to attenuate overinvestments designed to attract fans (the upgrading of the home stadium, the search for the best coach, and the coddling of players) college football teams share revenues within their conferences (Brown, 1994).[3]

Heintzelman, Salant, and Schott (2009) analyze the consequences of output sharing in an environment with negative externalities and unobservable effort. They consider a game where N players partition themselves into groups each of which shares output but not costs. They show that although forming N solo groups (no output sharing) generates too much effort due to negative externalities and forming one group with N members generates too little effort due to free riding, socially optimal effort can be achieved (or approximated due to an integer constraint) with output sharing groups of intermediate size.

Kamien, Muller, and Zang (1992) consider the effects of forming research joint ventures. They consider the case where every player invests in cost-reducing R&D and no benefits are shared, the analog of competing groups each with one member. They find that there is too much duplication. At the other extreme, they consider the case where every member of a research joint venture shares equally in the cost-reducing discoveries of the joint venture regardless of which member discovered it, the analog of one group containing all the researchers. They show that there is too little aggregate effort. They do not discuss the intermediate case, which would be socially optimal, where research joint ventures of the right size compete against each other.[4]

The strategic interactions in these models are similar to those in the Cournot model (1838). Indeed, "excessive" output in the Cournot model could *in theory* be reduced toward the monopoly level if competing subsets of players agreed to share within each group their gross revenues. This occurs at a local level within competing law partnerships and at an international level within ocean liner conferences which share gross revenues to facilitate cartelization (see Bennathan and Walters 1969).

[2]In Japan, one hundred forty-seven Japanese fisheries currently engage in output sharing. Platteau and Seki (2000) interviewed skippers in one such fishery. They concluded that "the desire to avoid the various costs of crowding while operating in attractive fishing spots appears as the main reason stated by Japanese fishermen for adopting pooling arrangements."

[3]Brown (1994) concludes that "The empirical results confirm the theoretical prediction that revenue sharing provides a disincentive to build a stronger team, other factors constant—-conferences which share more tend to have weaker teams. . . "

[4]For an alternative characterization of research joint ventures, see D'Aspremont and Jacquemin (1988) and Salant and Shaffer (1999).

Schott, Buckley, Mestelman, and Muller (2007) examine output sharing experimentally by exogenously dividing subjects into equal-size groups. They demonstrate that exogenous variations in group size affect subsequent behavior as predicted and that the appropriate exogenous group size results in socially optimal behavior.

In this paper, we investigate whether, given the opportunity to form equal-sized groups of different sizes, individuals will *choose* to form groups of the socially optimal size.[5] In addition, we investigate whether the output-sharing groups selected motivate subjects to invest efficiently. To do this, we conduct a laboratory experiment where subjects vote on the size of their output-sharing groups and then play an investment game in the chosen group structure. Finally, we also explore whether individuals choose the efficient group sizes and invest optimal amounts under different costs of investment. Establishing how players partition themselves endogenously is important, since, in the field, subjects will *choose* how many groups to form. If players turned out always to vote for a suboptimal number of groups, then our laboratory society would never reach efficiency even if, as in Schott, Buckley, Mestelman, and Muller (2007), it makes socially optimal choices when the optimal group structure is exogenously mandated.

In our experiments, individuals grouped into solo groups did overinvest relative to the socially optimal level. But as the group size increased, subjects invested smaller amounts. When given the opportunity to *choose* the size of their groups, most of the subjects voted for the group size that is socially optimal, and subjects cut the waste associated with the negative aggregate spillover problem on average by at least two-thirds in three of the cost treatments and by one-half in the remaining cost treatment. When we varied the opportunity cost of investment/effort, subjects tended to vote for the group size that became socially optimal given the new circumstance.

However, systematic departures from the theory were also noted. We find that *without exception* departures are in the direction of the socially optimal investment. Hence, as elsewhere in the literature, we find that laboratory behavior is more “cooperative” and “other-regarding” than a theory based on self-interested behavior would predict (i.e., Ostrom, Walker, and Gardner 1992; Ostrom, Walker, and Gardner 1994; Ledyard 1995; Camerer 2003; Falk, Fehr, and Fischbacher 2005).

The paper proceeds as follows: Section 2 describes our experimental design and procedures. Section 3 presents our theoretical hypotheses. Section 4 reports our experimental findings and the results of our hypothesis tests. Section 5 discusses directions for future research and concludes the

[5] Although ours is the first paper to investigate endogenous output-sharing groups in a setting with negative externalities, the formation of groups has been investigated in public goods environments (Page, Putterman, and Unel 2005; Ahn, Isaac, and Salmon 2009; Brekke, Hauge, Lind, and Nyborg 2009; Charness and Yang 2008).

paper.

3.2 The Experiment

3.2.1 Design and Procedures

We conducted 25 sessions, each with a different set of 6 participants. Most participants were undergraduates at [insert university name here]. Subjects earned experimental currency (tokens), which was converted at the conclusion of the session into US dollars (1 token = 0.01 US dollars). The experiment was programmed and conducted with the software z-Tree Fischbacher (2007). Sessions took approximately one hour and a half.

Each session was divided into six separate parts. Each of the first five parts (Parts I–V) consisted of a sequence of 5 rounds of decision making. Therefore, each subject went through 25 rounds in total. One aim of the first four parts was to give subjects experience investing as members of groups of different sizes. In Part V of the experiment, subjects *chose* the size of the groups endogenously. In Part VI, subjects completed a short questionnaire. At the end of the experiment, we randomly selected one round from each of the first five parts, added up the tokens each subject had earned in the selected rounds, converted that sum to dollars, added in the $5 show-up fee, and paid everyone. The average payment in the experiment was approximately $25 per subject.

In the first four parts, subjects were exogenously divided into groups of identical size: one-member groups, two-member groups, three-member groups, or a six-member group. Subjects were randomly rematched across groups in every round but played 5 consecutive rounds in each group size in order to gain experience. In total, there were 20 rounds in the first four parts. In order to control for order effects, the order of the first four parts was changed across sessions.

At the beginning of each decision round in the first four parts, participants were given 6 experimental tokens and had to decide how many of them $(0, 1, \ldots, 6)$ to invest in Project B. Whatever a subject did not invest in Project B was automatically invested in Project A. Denote x_{ik} as the investment in Project B by agent k in group i. Let Y_i^{-k} denote the aggregate investment in Project B by the other members of group i, X_{-i} denote the aggregate investment in Project B by other groups, and X denote the total investment in Project B by all 6 participants.

Project A had a fixed return of c tokens per token invested; i.e., the subject's earnings from Project A equaled c times his investment in Project A. Therefore, the "opportunity cost" of investing one additional token in Project B equaled c, the lost earnings from Project A.

The return per token invested in Project B, $P(X)$, was a decreasing linear function of the aggregate investment in Project B. Therefore, each token invested in Project B imposed a negative externality on everyone else–the essence of the *negative aggregate spillover* problem. For each token invested in Project B, the return from Project B was given by

$$P(X) = 200 - 5X.$$

Every member of a given group received an equal share of his or her group's return from Project B regardless of his or her own investment in that project. An individual's earnings from Project B (E_{ik}) depended on the participant's group investment in Project B and the group size (m):

$$E_{ik} = \frac{1}{m}(200 - 5X)(x_{ik} + Y_i^{-k}).$$

When there is no output sharing, an individual's earning is given by $E_i = (200 - 5X)x_i$. This functional form is consistent with the formulation that is found in the common-property literature Dasgupta and Heal (1985), the innovation-tournament literature (Baye and Hoppe, Theorem 1), and the sports-contest literature (Dietl et al. 2008, equation 2 with $\alpha = \gamma = 1$), among others.[6] In all these seemingly unrelated strategic interactions, the same game has been analyzed repeatedly in disparate literatures.[7]

Final earnings in each round (in tokens) were simply the sum of earnings from Project A and earnings from Project B:

$$\pi_{ik} = (6 - x_{ik})c + \frac{1}{m}(200 - 5X)(x_{ik} + Y_i^{-k}).$$

It can be seen that each individual pays the cost of his or her investment but shares the revenue equally with the members of his group.

In each of the five rounds of Part V, subjects first *voted* for one of the four group sizes. Then, subjects were divided up in groups of the size that won the most votes and played the investment

[6] The same model also appears in the "rent-seeking" literature (Chung 1996, equation 2). For other literatures where this model appears, see the excellent book-length survey Konrad (2009).

[7] In each case, player i chooses effort/investment x_i and achieves payoff $\frac{x_i}{x_i+X_{-i}}v(x_i + X_{-i}) - cx_i$ where $v(\cdot)$, the reward function, gives the value of the output or prize. When the payoff function of each player is rewritten as $x_i(P(x_i + X_{-i}) - c)$, where the strictly decreasing function $P(x_i + X_{-i}) = \frac{v(x_i+X_{-i})}{x_i+X_{-i}}$, the paternity of this ubiquitous model becomes apparent: it is the Cournot model (1838) in disguise (with x_i reinterpreted as effort instead of output). In the Cournot model, of course, the negative aggregate spillover problem results in larger industry output than a monopolist would choose. These literatures all assume that $v(\cdot)$ is strictly concave but differ in whether this function is single-peaked (as in the Cournot model) or strictly increasing (as in Baye and Hoppe 2003); the qualitative results in these literatures are unaffected by this minor difference in assumption. We chose the simpler of the two formulations as easier to explain to subjects. We assumed that $P(X)$ decreases linearly, which implies that $v(X)$ is single-peaked (a parabola).

game. In cases of a tied vote, the winner was chosen at random.[8]

In our experimental design, different group sizes are socially optimal under different treatments. In particular, as the opportunity cost of investing in Project B increases, the optimal group size decreases. Subjects in a given experimental session faced only one cost parameter and had to make investment decisions in all five parts of the experiment (25 rounds). A summary of the experimental design is provided in Table 3.1. As Table 3.1 reflects, the socially optimal group size is different for each treatment. For example, for opportunity cost $c = 20$, the optimal size of each group is 3 members (or, equivalently, the optimal number of groups is 2).

Table 3.1: Experimental Design

Cost parameter *c*	Efficient group size	Parts I – IV	Part V	Number of sessions	Number of subjects
1	6	Exogenous	Voting	5	30
20	3	Exogenous	Voting	5	30
55	2	Exogenous	Voting	5	30
100	1	Exogenous	Voting	5	30

Prior to the experiment, a test was administered to the subjects to make sure they understood the payoff consequences of their choices. The computer prevented anyone from beginning the session until *everyone* had a perfect score on the test.

During the experiment, subjects could either calculate their payoffs by hand or could utilize a "Situation Analyzer" provided to facilitate their calculations. A subject could enter his or her conjecture about (1) the total investment in Project B by others *inside* his or her group and (2) the total investment in Project B by subjects *outside* his or her group. The Situation Analyzer would then provide a table listing in one row the seven choices for investing in Project B $(0, 1, \ldots, 6$ tokens) and in the other row the total payoff from the two projects that the subject would earn if his or her two conjectures were accurate. Subjects were free to do such calculations by hand or to use the Situation Analyzer as often as they wanted before making a decision. The Situation Analyzer is shown in Figure 3.1.

To help subjects to make a decision, subjects were also reminded of their own investments, others' investments in their group, and the total investment, as well as their earnings from previous

[8] Note that in our experiment, voting is used to select group sizes and not some policy regarding the level of effort/investment. Voting is found to be useful as a tool for policy selection in common pool resources or public goods literature (Walker, Gardner, Herr, and Ostrom 2000; Tyran and Feld 2006; Putterman and Kamei 2010).

Figure 3.1: Situation Analyzer for Groups of Two

If the total investment by others outside your group in Project B is 12

If the total investment by others inside your group in Project B is 4

Analyze this situation

Your investment in B	0	1	2	3	4	5	6
Your earnings (A + B)	840.0	787.5	730.0	667.5	600.0	527.5	450.0

rounds.

After the session, we administered a short questionnaire. We asked subjects the basis of their investment decisions and the basis of their vote on group size. Responses clearly showed that subjects understood the experiment. Most of the subjects reported that they tried to maximize their monetary earnings.

3.3 Theoretical Predictions and Hypotheses

Theoretical predictions are based on Heintzelman, Salant, and Schott (2009). Given the size of each group and the opportunity cost of investing in Project B, every individual simultaneously invests to maximize his or her own payoff from the two projects. Proposition 3.1 summarizes the mean investment in Project B in the Nash equilibrium.

Proposition 3.1. *For a given opportunity cost (c) and group size (m), mean investment in Project B is* $\bar{x} = \frac{200-cm}{30+5m}$.[9]

We use this formula to calculate the mean equilibrium investment in Table 3.2 for each of the 16 group size-cost combinations in the various experimental treatments. For any opportunity cost (c), note that Nash equilibrium investment in Project B decreases with the size of each group and that, for any group size (m), investment also decreases with the opportunity cost.

Proposition 3.2 summarizes the mean investment in Project B that would maximize the aggregate revenue from the two projects. We refer to this as the "socially efficient investment level" and denote it x^*. It is simply 1/6 of the aggregate investment that maximizes $XP(X) + (36 - X)c$.

[9] Total investment within each group is uniquely determined in the equilibrium, but not the investment of individual members of a group. Therefore, we focus on the mean investment level. See Heintzelman, Salant, and Schott (2009) for more details.

Proposition 3.2. *To maximize social surplus, mean investment in Project B must be* $x^* = \frac{200-c}{60}$.

This formula is used to calculate the socially efficient levels in the last row of Table 3.2. Note that as the opportunity cost rises, the socially efficient investment level falls.

Equilibrium investment in Project B exceeds the socially optimal level when no one shares (one-person groups) and falls short of the socially optimal level when everyone shares (six-person group). Moreover, as the group size increases from one to six, aggregate investment in Project B declines monotonically. There is, therefore, a unique real number m (not necessarily an integer) which induces mean investment in the Nash equilibrium ($\bar{x}$) to equal socially efficient investment (x^*). When the integer constraint is respected, one of the six group sizes will generate a larger social surplus in the Nash equilibrium than any other group size and is predicted to increase efficiency close to the socially optimal level. We refer to this group size as the "partnership solution" (Heintzelman et al. (2009)). Note that the partnership solution is a self-enforcing mechanism that requires neither monitoring of individual behavior nor intervention of the government.

In Table 3.2, in every cost column there is an entry in bold-face. The associated row is the partnership solution for that particular opportunity cost. The entry in bold is the predicted mean investment in the partnership solution for that opportunity cost.[10]

Table 3.2: Predicted Mean Investment in Project B

Group size	Cost = 1	Cost = 20	Cost = 55	Cost = 100
1	5.69	5.14	4.14	**2.86**
2	4.95	4	**2.25**	0
3	4.38	**3.11**	0.78	0
6	**3.23**	1.33	0	0
	Socially efficient level = 3.32	Socially efficient level = 3	Socially efficient level = 2.42	Socially efficient level = 1.67

Proposition 3.3 follows from the definition of the partnership solution:

Proposition 3.3. *For a given opportunity cost (c), every individual strictly prefers the partnership solution to any other group size.*

[10] In this experiment, subjects faced a discrete action space. Though the theoretical predictions were generated from a game with continuous actions, the assumption of discrete actions does not change the predictions. More specifically, suppose agents choose a noninteger investment level x for Project B in the symmetric equilibrium of the continuous investment game. Then, in the discrete version, there is an equilibrium in which every player chooses the integer above x or below x, or mixes between the two. As a result, both the actions and the payoffs in the discrete and continuous versions are very similar.

The partnership solution is the Condorcet winner, because it would be chosen unanimously if every subject voted for the group size he most preferred. Of course, as in many voting games, inferior alternatives can also be supported as Nash equilibria since if everyone expects that everyone else is voting for the same alternative, then no one can change the outcome by deviating unilaterally. However, these spurious Nash equilibria can be eliminated by iterative elimination of weakly dominated strategies. The partnership solution would then be predicted to receive the most votes.[11]

In this paper, we test the following hypotheses as well as the point predictions presented in Table 3.2:

Hypothesis III.1. *For a given opportunity cost (c) of investing in Project B, mean investment in that project strictly decreases with the size of the groups.*

Hypothesis III.2. *Mean investment in Project B decreases with the opportunity cost of investing in that project for a given group size.*

Hypothesis III.3. *Subjects will "elect efficiency"—they will vote for the group size predicted to generate the largest social surplus in the Nash equilibrium.*

The experimental data and findings are presented in the next section.

3.4 Data Analysis

3.4.1 Exogenous Groups and Investment Decisions

Figure 3.2 shows the average investment corresponding to each opportunity cost parameter in the first 20 rounds (Parts I–IV). For simplicity, group sizes are presented in the following order: one-member, two-member, three-member, and six-member groups, although orders were randomized during the sessions.[12] Consistent with the theoretical predictions, contributions decrease with the group size for any cost level.

Theoretical predictions and the observed mean levels of investment in Project B are provided in Table 3.3. Theory predicts that the socially optimal group size decreases with cost. Observed mean investment and predicted investment in Project B are shaded for the theoretically optimal

[11]We piloted a second voting mechanism which needs no such refinement since its Nash equilibrium is unique: after each subject had voted in the pilot, one of the six subjects was randomly chosen to be "dictator," and his or her vote determined the partnership structure. Since every subject anticipated being chosen as the dictator with positive probability, each subject should have been motivated to vote for his or her most preferred alternative. However, we were unable to distinguish behavior under the two voting schemes. Hence, we used the more familiar nondictatorial scheme for this paper.

[12]We do not observe any order effects in our data.

Figure 3.2: Mean Investment

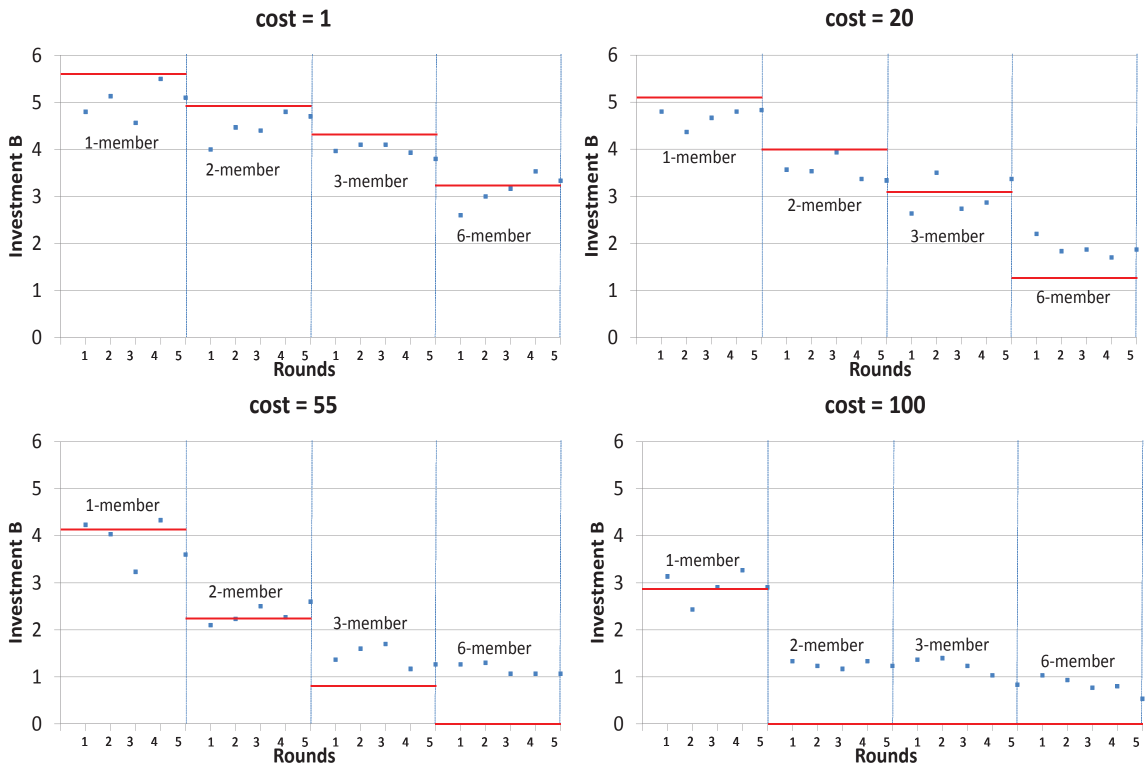

group sizes. Observed mean investment at the optimal group size for each cost is surprisingly close to the theoretical predictions and the socially optimal level of investment.

We performed some nonparametric tests by using independent observations (one data point per session). One-sided sign tests confirm that there are no significant differences between the observed levels of investment and the theoretical predictions at the optimal group sizes ($p - values > 0.1$). For nonoptimal group sizes, point predictions do not hold in general (p-values are generally less than 0.05).[13] However, all deviations are toward the socially optimal level.

Result III.1. *Theoretical predictions hold at the optimal group sizes. However, there are deviations from quantitative predictions for other group sizes. When data are not consistent with the predicted levels, deviations are in the direction of socially optimal level in all cases.*

Table 3.4 shows the observed mean payoff for each cost and group size. For cost levels $c = \{1, 20, 55\}$ theoretically predicted optimal group size generates the highest level of payoff. Note that for $c = 100$ theoretically predicted optimal group size is 1 (no output sharing). However, for $c = 100$, higher levels of payoff are achieved with group sizes more than 1. One possible explanation

[13] The two exceptions are when cost is 20 and group size is six and when cost is 55 and group size is one. In these cases, investments are not significantly different than the predicted levels.

Table 3.3: Predicted versus Observed Mean Investment

	c = 1		c = 20		c = 55		c = 100	
Group size	Predicted	Observed	Predicted	Observed	Predicted	Observed	Predicted	Observed
1	5.69	5.02 (1.36)	5.14	4.69 (1.55)	4.14	3.89 (1.50)	2.86	2.93 (1.70)
2	4.95	4.47 (1.41)	4	3.55 (1.54)	2.25	2.34 (1.68)	0	1.26 (1.38)
3	4.38	3.98 (1.55)	3.11	3.02 (1.56)	0.78	1.42 (1.35)	0	1.17 (1.46)
6	3.23	3.13 (1.85)	1.33	1.89 (1.59)	0	1.15 (1.47)	0	0.81 (1.12)
	Socially efficient investment = 3.32		Socially efficient investment = 3		Socially efficient investment = 2.42		Socially efficient investment = 1.67	
Standard deviations are in parentheses Number of observations = 150 per cell								

is that, as Table 3.3 shows, theoretically predicted level of investment is not very close to the socially efficient level (since it is not possible to divide individuals into noninteger group sizes). Even though the mean investment with solo groups is not significantly different than predicted, the deviations we observe in the other group sizes affect the payoffs in an unpredicted way.[14]

Table 3.4: Predicted versus Observed Mean Payoff

	c = 1		c = 20		c = 55		c = 100	
Group size	Predicted payoff	Obs. ave. payoff	Predicted payoff	Obs. ave. payoff	Predicted payoff	Obs. ave. payoff	Predicted payoff	Obs. ave. payoff
1	168	241 (88.03)	252	295 (90.61)	416	427 (75.65)	640	625 (61.61)
2	256	286 (75.70)	360	370 (67.48)	504	497 (101.21)	600	670 (109.19)
3	302	314 (58.25)	390	377 (64.86)	425	465 (85.11)	600	671 (132.10)
6	336	323 (22.34)	307	341 (48.99)	330	443 (88.00)	600	656 (105.90)
	Socially efficient payoff = 336		Socially efficient payoff = 390		Socially efficient payoff = 505		Socially efficient payoff = 683	
Standard deviations are in parentheses Number of observations = 150 per cell								

For cost levels $c = \{1, 20, 55\}$, we test whether the Partnership Solution improves the payoff of participants relative to the case where there is no output sharing (being solo). By using matched-pair sign-rank tests, we confirm that the Partnership Solution increases the payoffs. In particular, we compare the mean payoff levels at the socially optimal group size with the mean payoff levels at the group size of one. Each individual's payoff increases with the Partnership Solution and the difference is significant at the 5% level.

[14] For group sizes greater than 1, complete free riding is not observed as predicted. This is consistent with behavior observed in public goods experiments. It has been documented that subjects do not free ride completely (see Ledyard 1995).

For $c = 100$, the group size of 1 brings the lowest payoff, even though it was the theoretically optimal group size (p-values for all pairwise comparisons are 0.04). Output sharing seems to help individuals even in situations where theoretically it is not the case.

Result III.2. *Output sharing improves payoffs when groups are exogenously formed.*

We complement nonparametric tests with a regression analysis. We investigate the impact of different group sizes, costs, the order of presenting group sizes, and rounds on individual investment decisions by running ordinary least squares estimation with robust standard errors (see Table 3.5).[15]

Regression results (specification 1) show that for a given cost level, an increase in the group size decreases the level of investment in Project B. We also see that there is a negative relationship between the level of investment and the opportunity cost parameter, c. Specifications 2–4 show that these results continue to hold even when we add control variables or when we include the different treatments as dummy variables.[16] In addition, we see that the order of treatments and experience do not affect investment decisions.[17] In summary, one cannot reject hypotheses III.1 and III.2. Our results are robust to different estimation methods.[18]

Result III.3. *The data are consistent with the (qualitative) theoretical predictions. For each cost level, investment decreases with group size. Moreover, investment decreases with cost for a given group size.*

3.4.2 Voting for Group Size: The Plurality Rule

Table 3.6 presents the percentage of votes that each group size received for each cost level. There are 150 observations for a given level of cost and group size. Except for $c = 100$, groups frequently vote for the theoretically predicted optimal group size. Approximately 60% of the votes are socially optimal for $c = \{1, 55\}$, and approximately 40% of the votes are socially optimal for $c = 20$.

For each cost parameter, we test whether one can reject the null hypothesis that the proportion of votes is 25% for each group size. For $c = \{1, 20, 55\}$, one can strongly reject this null hypothesis (chi-square goodness of fit test, $p - values < 0.01$). For $c = 100$, one cannot reject that the proportion of votes is 25% for each group size ($p - value = 0.10$). More important, the highest

[15] Data are clustered by 20 sessions.

[16] We find that the coefficient of grsize2 is significantly smaller than the coefficient of grsize3, and the coefficient of grsize3 is significantly smaller than the coefficient of grsize6 ($p - values < 0.01$). We find the same result for cost parameters as well.

[17] Note that the variable round takes values $1, 2, ..., 5$.

[18] For robustness checks, we have also conducted fixed-effect regressions both at the individual and at the session levels. Group size affects investment negatively for all cost levels. In addition, round seems to have a small but significantly negative effect for cost levels greater than 1. Results are available upon request.

Table 3.5: Ordinary Least Squares Results

Dependent var: Investment B	1	2	3	4
groupsize	−0.42**	−0.42**	−0.43**	
	(0.02)	(0.02)	(0.02)	
cost	−0.03**	−0.03**	−0.03**	
	(0.00)	(0.00)	(0.00)	
round		0.00	0.00	
		(0.02)	(0.02)	
grsize2				−1.23**
				(0.12)
grsize3				−1.74**
				(0.14)
grsize6				−2.39**
				(0.13)
cost20				−0.86**
				(0.12)
cost55				−1.95**
				(0.11)
cost100				−2.61**
				(0.09)
round2				0.04
				(0.10)
round3				−0.03
				(0.09)
round4				0.09
				(0.10)
round5				−0.00
				(0.10)
phase2			−0.10	−0.00
			(0.17)	(0.10)
phase3			−0.01	0.08
			(0.18)	(0.12)
phase4			−0.08	0.02
			(0.23)	(0.14)
Constant	5.20**	5.19**	5.25**	5.45**
	(0.10)	(0.09)	(0.16)	(0.16)
Observations	2,400	2,400	2,400	2,400
R-squared	0.384	0.384	0.384	0.431
Robust standard errors in parentheses ** p < 0.01, * p < 0.05				

percentage of votes is for the socially optimal group sizes. In particular, for $c = 1$, group size 6 received the highest number of votes; for $c = 20$, group sizes 2 and 3 received the highest number of votes; and for $c = 55$, group size 2 received the highest number of votes (proportion tests, $p - values < 0.01$). For $c = 100$, group size 3 received significantly more votes than the socially optimal level of one (proportion test, $p - value = 0.049$).

Result III.4. *For $c = \{1, 20, 55\}$, the highest proportion of votes is received by the corresponding*

Table 3.6: Percentage of Votes in Part V

Group size	c = 1	c = 20	c = 55	c = 100
1	8	12.7	7.3	22
2	16	39.3	57.3	24.7
3	16.7	39.3	11.3	33.3
6	59.3	8.7	24	20

socially optimal group sizes. (This holds weakly for $c = 20$*.)*

Result III.4 shows that participants choose to form output-sharing groups for all cost levels. In addition, we conduct a multinomial logit regression analysis to test whether votes are affected by cost, previous earnings and experience.[19] We construct a new variable, bestgroup, which takes value 1, 2, 3 if a subject earned the most money in Parts I–IV when the group size is 1, 2, 3, respectively, and takes value 4 if a subject earned the most money when the group size is 6.[20] Regressors are jointly significant at the 0.05 level (Wald chi-square = 74.93, $p - value < 0.01$). In addition, we find that both cost and bestgroup significantly affect votes (Wald tests, $p - value = 0.03$ and $p - value < 0.01$ respectively). However, coefficient estimates of round are not jointly statistically significant ($p - value = 0.40$). Table 3.7 presents the marginal effects after a multinomial logit regression. Robust standard errors are provided in parentheses.

Table 3.7: Multinomial Logit Regression – Marginal Effects

	Dependent variable = vote			
VARIABLES	Group size = 1	Group size = 2	Group size = 3	Group size = 6
cost	0.001	–0.000	0.001	–0.002*
	(0.001)	(0.001)	(0.001)	(0.001)
bestgroup	–0.037	–0.266**	0.059	0.245**
	(0.034)	(0.052)	(0.036)	(0.032)
round	0.007	0.005	–0.023	0.011
	(0.011)	(0.012)	(0.013)	(0.017)
Observations	600	600	600	600
Robust standard errors in parentheses ** p < 0.01, * p < 0.05				

We see that the probability of voting for group size 6 significantly decreases with cost and increases with bestgroup, whereas the probability of voting for group size 2 decreases with bestgroup.[21] These findings are consistent with the theoretical predictions. A simple correlation analysis

[19] Since utilities from different group sizes do not need to be ordered, a multinomial logit regression analysis is more suitable than an ordered logit regression analysis. In addition, we have performed OLS regressions, and qualitative results did not change.

[20] The earnings in each part are calculated by adding up each payoff from the 5 corresponding rounds.

[21] Since bestgroup is a discrete variable, we have also looked at the predicted probabilities for each group size under

also confirms that votes are negatively correlated with cost (–0.21) and positively correlated with bestgroup (0.46).

Result III.5. *Votes are affected by both the cost parameter and the previous earnings at different group sizes. Votes do not change significantly as subjects get more experienced with voting.*

Table 3.8: Mean Investment and Payoff Conditional on Chosen Group Size

c	Group size	Frequency (out of 25)	Investment		Payoff	
			Predicted	Observed	Predicted	Observed
c = 1	1	2	5.69	5.75 (0.62)	168	158 (21.45)
	2	2	4.95	4.67 (1.23)	256	278 (53.76)
	3	1	4.38	4.50 (1.22)	302	294 (58.43)
	6	20	3.23	3.26 (1.71)	336	323 (20.82)
			Predicted	Observed	Predicted	Observed
c = 20	1	3	5.14	4.89 (1.49)	252	266 (102.95)
	2	12	4.00	3.74 (1.65)	360	369 (86.64)
	3	10	3.11	3.00 (1.28)	390	383 (48.70)
	6	0	1.33	–	307	–
			Predicted	Observed	Predicted	Observed
c = 55	1	0	4.14	–	416	–
	2	18	2.25	2.20 (1.37)	504	498 (88.20)
	3	1	0.78	0.83 (0.98)	425	430 (57.47)
	6	6	0	1.11 (1.69)	330	442 (95.41)
			Predicted	Observed	Predicted	Observed
c = 100	1	7	2.86	2.67 (1.51)	641	648 (39.62)
	2	6	0	1.00 (1.39)	600	659 (111.82)
	3	11	0	0.85 (1.18)	600	661 (105.97)
	6	1	0	0.50 (0.83)	600	643 (83.67)

Table 3.8 presents the voting outcomes, mean investment decisions and payoffs conditional on the chosen group size.[22] As in the exogenous groups, we see that participants choose investment levels that are consistent with the theoretical predictions at the socially optimal group sizes (all

each possible value of bestgroup. We have observed similar results.

[22]Since ties are broken randomly, even though there are equal number of votes for group sizes 2 and 3 when cost is 20, group size 2 won the voting more frequently than group size 3.

p-values are greater than 0.27).[23] Moreover, qualitative results are similar to the case when groups are exogenously imposed: investment decreases with the group size ($p - value < 0.01$) and cost ($p - value < 0.01$). Regression results are available from the authors.

Result III.6. *Mean investment levels in Part V are not significantly different than theoretically predicted levels at the socially optimal group sizes. In addition, investments are consistent with the (qualitative) theoretical predictions. Investment decreases with group size and cost.*

Finally, we compare the efficiency of endogenous group formation with the case of exogenous groups. Efficiency of each part is defined by the observed average payoff divided by socially optimal payoff. In Table 3.9, we provide the efficiency levels in all parts for each cost treatment. As expected, efficiency levels are quite large. Endogenous group formation increases efficiency compared with the case of no output sharing for all cost levels. In particular, efficiency loss decreased by 50% for $cost = 100$ and by 68% to 71% for the other cost levels.

Table 3.9: A Comparison of Efficiency Levels

	Group size	c = 1	c = 20	c = 55	c = 100
Exogenous	1	0.72	0.76	0.84	0.92
	2	0.85	0.95	0.98	0.98
	3	0.93	0.97	0.92	0.98
	6	0.96	0.87	0.88	0.96
Endogenous	voting	0.91	0.93	0.95	0.96

3.5 Discussion and Conclusions

When every individual's effort imposes negative externalities on his competitors, competition results in excessive aggregate effort. This explains overfishing when competing on common properties, excessive proposal polishing in grant competitions, duplication in innovation tournaments, and excessive talent searches among competing sports teams. In theory, one way to curb these excesses is for subsets of competitors to form groups which share output or revenue. If the right number of groups forms, Nash equilibrium aggregate effort should fall to the socially optimal level.

We investigated experimentally whether individuals *in fact* manage to form the efficient number of groups and to invest within the chosen groups as theory predicts. We find that output sharing attenuates the negative aggregate spillover problem independent of the opportunity cost of investing. Consistent with theoretical predictions, we find a negative relationship between the aggregate

[23] We focus on the socially optimal group size, since votes are more often for the optimal group size. Therefore, there are not too much data available on the other group sizes. In fact, there are too few data points for many of the nonoptimal group sizes, which makes statistical testing not very meaningful.

investment levels (the counterpart to fishing effort, R&D effort, etc.) and group size. For a given group size, we show that aggregate investment decreases as the opportunity cost of investing in it increases. More importantly, we show that socially optimal group sizes are the most common outcome of the endogenous group formation stage under most of the cost parameters.

Regarding the point predictions, we find that partnership solution (exogenous implementation of socially optimal group size) generates theoretically predicted levels of investment. However, in general, theory does not predict the magnitudes very well for the nonoptimal group sizes. For any deviations from equilibrium predictions, we see that investments shift toward the efficient outcome.[24] One explanation for this is that individuals are altruistic. If individuals care not only for themselves but also for others, then one would expect to see higher levels of efficiency than a theory predicated on the assumption of self-interested behavior would predict (except when the theory predicts socially optimal outcomes). This is highly consistent with our experimental data. Moreover, this type of behavior has been commonly observed in other experimental studies on common-pool resources and public goods (see Ostrom, Walker, and Gardner 1994; Ledyard 1995).

Future research should address the stability of the partnership mechanism and its sensitivity to inter-subject communication. By stability, we mean migrations of subjects among existing groups or from an existing group to a newly formed group.[25] The effect of inter-subject communication on the Partnership Solution is the subject of a recent study by Buckley, Mestelman, Muller, Schott, and Zhang (2009). They find that when individuals within the same output-sharing group are able to communicate, free riding decreases. It is unclear from their work whether similar results would occur if subjects collectively *chose* their group size; moreover, communication may affect the choice of group size itself. We leave the investigation of such interplay between communication and endogeneity to future research.

[24] We focus on the exogenous groups since there are very few data points for statistical testing for endogenous groups at the nonoptimal group sizes.

[25] After the voting stage but before the investment stage, a migration stage could be inserted. Subjects could be permitted to migrate simultaneously from the group to which they are assigned after the voting stage to another group of their choosing. Heintzelman, Salant, and Schott (2009) predict that no such migrations should occur. However, they also predict that migrations to newly formed solo partnerships would occur *unless* there was a direct cost to such migrations or there was a benefit to team production which would be lost by going solo. We plan to test these predictions in future work.

CHAPTER IV

Omitted Proofs

4.1 Any PCE Payoff is an MPCE Payoff: Proof of Lemma I.1

Let $W \subset \mathbb{R}^n$ be a compact, convex set with extreme points denoted ext W. The continuation value function $\Bbbk : A^2 \to W \subset \mathbb{R}^n$ has the *bang-bang* property if $\Bbbk(a, \tilde{a}) \in \text{ext } W$ for all action profiles $a \in A$ and recommendation profiles $\tilde{a} \in A$.

We first argue that any continuation value function can be replaced with one that takes values in ext W.

Claim IV.1 (Bang-Bang)**.** *Any continuation value function is equivalent to one with the bang-bang property.*

Proof of Claim IV.1: We adapt the proof of Theorem 3 in APS, accounting for correlation and a finite domain of the continuation value function.[1] For a bounded set $W \subset \mathbb{R}^n$, let $\mathcal{K}(W)$ be the set of all functions from $A \times A$ to W, and $\mathcal{K}(W|w) \subseteq \mathcal{K}(W)$ the set of continuation value functions that support w on W. Since $\mathcal{K}(W) = W^{|A|^2}$, and W is compact, it is compact in the product topology, by Tychonov's Theorem. Next, since a convex combination of admissible pairs is also an admissible pair , $\mathcal{K}(W|w)$ is a convex set. As a closed subset of a compact set, it is compact. By the Krein-Milman Theorem, any $\Bbbk \in \mathcal{K}(W|w)$ can be written as a convex combination of extreme points of $\mathcal{K}(W|w)$. Finally, linearity of incentives and payoffs implies that $\hat{\Bbbk}$ is a convex combination of extreme points of $\mathcal{K}(\text{ext } W|w)$, and consequently has the bang-bang property. □

Proof of Lemma I.1: Let $[x_1, \ldots, x_m]$ be the convex hull of the points $(x_1, \ldots, x_m)$. Let V_{PCE} be the set of PCE payoffs. Fix a PCE $\lambda \in \Lambda$ with payoff $w \in V_{PCE}$. Define the product space $\mathcal{V} \equiv$

[1] In APS, an equilibrium prescribes continuation behavior for each of a continuum of possible public signals. This required an appeal to Aumann (1965) for technical reasons. In our context, a continuation value function is defined on a finite set. The set of continuation value functions, therefore, is a simpler object that can be treated with simpler mathematical tools.

${V_{PCE}}^{|A|^2}$. To prove that the payoff w is attainable in an MPCE, we show that there exists a correlation device $\lambda_M \in \Lambda_M$ that delivers the payoff w and is incentive compatible. Thus, we want to show that the payoff vector w is supported by the convex hull of a self-generating set of $|A|^2$ payoff vectors.

Any continuation value function can be written as an ordered $|A|^2$-tuple of payoff vectors, one for each action profile and recommendation. Define the correspondence on $|A|^2$-tuples of payoff vectors $\phi : \mathcal{V} \to 2^{\mathcal{V}}$ by

$$\phi(v_1, \dots, v_{|A|^2}) = [\mathcal{K}(V_{PCE}|v_1), \dots, \mathcal{K}(V_{PCE}|v_{|A|^2}))] \cap [v_1, \dots, v_{|A|^2}]$$

The correspondence ϕ maps $|A|^2$-tuples of payoff vectors to the convex hull of supporting sets of $|A|^2$-tuples of payoff vectors.

We now claim that the correspondence ϕ satisfies the hypotheses of the Kakutani Fixed Point Theorem. Since V_{PCE} is non-empty, compact and convex, $\mathcal{V}$ is non-empty, compact and convex. The correspondence has non-empty values: since v_j is a PCE payoff, $\mathcal{K}(V_{PCE}, v_j)$ is not empty. Since by Claim IV.1 continuation payoffs can equivalently be taken from ext V_{PCE}, the intersection with the convex hull of an arbitrary set of PCE payoffs is non-empty. Furthermore, ϕ takes compact convex values as the intersection of two compact, convex sets. By Claim IV.1 and the Theorem of the Maximum, $\mathcal{K}(V_{PCE}, v_j)$ is upper hemi-continuous in v_j. Similarly, $[v_1, \dots, v_{|A|^2}]$ is upper hemi-continuous. Then ϕ is the intersection of upper hemi-continuous correspondences and therefore also upper hemi-continuous. Thus, by the Kakutani Fixed Point Theorem there exists a fixed point $(v_1^*, \dots, v_{|A|^2}^*)$.

For each element v_j^* of the fixed point, $j = 1, \dots, |A|^2$, there exists a probability distribution μ_j^* on A^2 used to enforce it, since each is a PCE payoff. Then the device λ_M making recommendations according to μ_j^* for $j = 1, \dots, |A|^2$ is Markovian and incentive compatible by construction. □

4.2 Characterization of MPCE: Proof of Theorem I.1

Part (a) Factorization: First we show that if W is self-generating, then $B(W) \subseteq V^*$. For any payoff vector $w \in B(W)$ there exists a pair $(\mu, \Bbbk)$ that enforces w on W. Since W is self-generating, $\Bbbk(a, \tilde{a}) \in W$ for all outcomes $(a, \tilde{a})$. Each payoff $\Bbbk(a, \tilde{a})$ is enforced on W. In this way, we can (by the Axiom of Choice) recursively define a PCE by constructing admissible pairs ad infinitum. By Lemma I.1, the PCE payoff w is an MPCE payoff. Thus, $W \subseteq V^*$. Next, we prove that V^* is a fixed point of $B(\cdot)$. Since V^* contains every self-generating set, we need only show that V^* is self-generating. Consider an MPCE payoff $w \in V^*$. There exists a pair $(\mu, \Bbbk)$ such that $\Bbbk(a, \tilde{a}) \in V^*$ for each pair of action and recommendation profiles $(a, \tilde{a})$. Hence, w is admissible w.r.t. V^*, or equivalently that $w \in B(V^*)$.

Finally, suppose that there exists a fixed point W of $B(\cdot)$ that strictly contains V^*. Then W is self-generating, and so is contained in the MPCE set V^*. This contradicts the premise that W strictly contains V^*. So V^* is the largest fixed point of $B(\cdot)$. □

Part (b) Compact and Convex: First, we want to show that $B(W)$ is compact if W is compact. Since $B(W)$ is bounded, by the Heine-Borel Theorem it is compact if it is also closed. Consider a sequence $\{b_j\}$ in $B(W)$ that converges to some $b \in \mathbb{R}^n$. Each $b_j \in B(W)$ is supported on W by an admissible pair $(\mu_j, \Bbbk_j)$. Endow the space of such functions that map $A \times A^2$ into $\Delta(A) \times W$ with the weak-* topology (i.e. pointwise convergence). The sequence is bounded, and so by the Bolzano-Weierstrass Theorem it has a convergent subsequence $\{\mu_l, \Bbbk_l\}$. The weak inequalities that define incentives are satisfied pointwise in the sequence $\{\mu_l, \Bbbk_l\}$, and hence are also by the limit $(\mu, \Bbbk)$, which thus enforces $b \in \mathbb{R}^n$. Then $b \in B(W)$, and so $B(W)$ is closed. □

Part (c) Contains SPE Payoffs: Since the mediated game has perfect monitoring of actions, players may ignore the correlation device, and instead play the subgame perfect equilibrium behavior after every history. □

Part (d) Nondecreasing δ: The proof is very similar to that of APS, Theorem 6. □

4.3 Algorithm: Proof of Theorem I.2

We extend the methods of Judd, Yeltekin, and Conklin (2003) to allow for correlation. Let $\mathbb{W}$ be the set of all convex subsets of V, partially ordered by set inclusion. Then the operator $B(\cdot)$ is monotone on the complete lattice $\mathbb{W}$. By Tarski's Fixed Point Theorem, $B(\cdot)$ has a largest fixed point V^*. Let $W^0 = V$ and recursively define $W^k = B(W^{k-1})$ for $k = 1, 2, \ldots$. First, by monotonicity $V^* = B(V^*) \subseteq B(W^0) = W^1$. Next, suppose that $V^* \subseteq W^k$. Monotonicity again yields $V^* = B(V^*) \subseteq B(W^k) = W^{k+1}$. By induction, $V^* \subseteq W^k$ for all $k = 1, 2, \ldots,$

The sequence $\{W^k\}_{k=0}^{\infty}$ is bounded and monotone, and therefore converges (in the Hausdorff topology) to a point in the complete lattice $\mathbb{W}$. Let $W^\infty = \lim_{k\to\infty} W^k$. This limit is a fixed point of $B(\cdot)$, and by construction contains V^*. But V^* cannot be a strict subset of W^∞, since that would imply that V^* is not the largest fixed point of $B(\cdot)$, contrary to Theorem I.1. □

4.4 MPCE as an Upper Bound: Proof of Theorem I.3

At the information set h_i^t, player i believes that the other players' private history profile is h_{-i}^t with posterior probability $\mu_{i,s}^t(h_{-i}^t|h_i^t)$, and that their period t action profile is a_{-i} with posterior probability

$$\beta_i^t(a_{-i}|h_i^t, s) = \sum_{h_{-i}^t \in H_{-i}^t} \mu_i^t(h_{-i}^t|h_i^t, s)\ s_{-i}(a_{-i}|h_{-i}^t)$$

Player i's *continuation payoff* under the strategy profile s at the private history h_i^t is therefore

$$\kappa_i^t(h_i^t|s) = (1-\delta)E\left[\sum_{r=t+1}^{\infty} \delta^{r-t-1} u_i(\beta_i^r)\ \middle|\, h_i^t, s\right] \tag{4.1}$$

where $u_i(\beta_i^t|h_i^t, s) = \sum_{a_{-i}\in A_{-i}} u_i(s_i(h_i^t), a_{-i})\beta_i^t(a_{-i}|h_i^t, s)$. Then player i's expected payoff under the strategy profile s at the private history h_i^t is

$$v_i^t(s|h_i^t) = (1-\delta)u_i(\beta_i^t)|h_i^t, s) + \delta\kappa_i^t(h_i^t|s)$$

As is well-known, a strategy profile s is a sequential equilibrium if and only if there are no profitable one-shot deviations. This is equivalent to

$$(1-\delta)u_i(\beta_i^t|h_i^t, s) + \delta\kappa_i^t(h_i^t|s) \geq (1-\delta)u_i(\beta_i^t|h_i^t, \tilde{s}_i, s_{-i}) + \delta\kappa_i(h_i^t|\tilde{s}_i, s_{-i}) \tag{4.2}$$

for all players i, private histories h_i^t, and strategies $\tilde{s}_i \neq s_i$.

Recall that s and v denote, respectively, the strategy profiles and payoffs in $G_\psi(\delta)$, and $\mathbb{s}$ and $\mathbb{v}$ denote, respectively, the strategy profiles and payoffs in $G^\lambda(\delta)$.

Claim IV.2 (The Correlation Device)**.** *For any strategy profile $s \in \mathcal{S}$ of $G_\psi(\delta)$, there exists a correlation device $\lambda_s \in \Lambda$ and strategy $\mathbb{s} \in \mathbb{S}$ in the mediated game that induces the same outcome in $G^{\lambda_s}(\delta)$ as s does in $G_\psi(\delta)$.*

Proof of Claim IV.2: For any strategy profile $s \in \mathcal{S}$, let $\beta^t\left(a^t|(a^1, \ldots, a^{t-1}), s\right)$ be the induced posterior probability of the action profile a^t in period t given the action history $(a^1, \ldots, a^{t-1})$. The action mixture in period 1 is simply $\beta^1(a^1) = \alpha^1(a)$. Given the realized action profile a^1, action profile a^2 occurs with chance $\beta^2(a^2|a^1) = \sum_{m^1\in M} \psi(m^1|a^1)s(a^2|a^1, m^1)$ using the joint density of signals $\psi(\cdot|a^1)$. In general,

$$\beta^t\left(a^t|s,(a^1,\ldots,a^{t-1})\right) = \sum_{(m^1,\ldots,m^{t-1})\in M^{t-1}} s\left(a^t|(a^1,\ldots,a^{t-1}),(m^1,\ldots,m^{t-1})\right)\prod_{k=1}^{t-1}\psi(m^k|a^k)$$

For all action histories $\mathbb{h}^t \in \mathbb{H}^t$, define $\lambda_s(\mathbb{h}^t) = \beta^t(\cdot|s,\mathbb{h}^t)$. Then, the recommendation distribution of λ_s coincides with the distribution of actions in $G_\psi(\delta)$. Call $\bar{\mathbb{s}}$ the *obedient strategy* in $G^{\lambda_s}(\delta)$ — namely, where players obey the recommendation of the correlation device λ after all histories. Since λ_s recommends the same outcome as w, the obedient strategy $\bar{\mathbb{s}}$ in $G^{\lambda_s}(\delta)$ delivers the same outcome as s. □

We must prove that obeying λ_s is a mutual best response, or $\mathbb{v}_i^t(\bar{\mathbb{s}}|\lambda_s) \geq \mathbb{v}_i^t(\mathbb{s}'_i,\bar{\mathbb{s}}_{-i}|\lambda_s)\ \forall \mathbb{s}'_i \in \mathbb{S}$. We'll argue that for every deviation $\mathbb{s}'_i \in \mathbb{S}_i$ in the mediated game, there is a corresponding strategy $s'_i \in \mathcal{S}_i$ with $\mathbb{v}_i^t(\mathbb{s}'_i,\bar{\mathbb{s}}_{-i}|\lambda_s) = v_i^t(s'_i,s_{-i})$. Namely, any deviation in the meditated game yields the same payoff as some strategy in the repeated game of private monitoring; this cannot be a profitable deviation against the sequential equilibrium profile s_{-i}. So $\mathbb{v}_i^t(\bar{\mathbb{s}}|\lambda_s) = v_i^t(s) \geq v_i^t(s'_i,s_{-i}) = \mathbb{v}_i^t(\mathbb{s}'_i,\bar{\mathbb{s}}|\lambda_s)$, as required.

Claim IV.3 (Verifying Incentives)**.** *If $s \in \mathcal{S}$ is a sequential equilibrium strategy of $G_\psi(\delta)$, then the correlation device $\lambda_s \in \Lambda$ is a PCE of $G(\delta)$.*

Proof of Claim IV.3: By the one-shot deviation principle, the obedient strategy is a best reply to itself iff there is no history after which a player would choose to disobey his recommendation once, and return to the obedient strategy thereafter. So, it suffices to restrict attention to alternative strategies that differ from the obedient strategy in one history. Consider a history $\mathbb{h}^t \in \mathbb{H}^t$ at which strategy $\mathbb{s}'_i$ instead plays the action a'_i in period t. Let $H(\mathbb{h}^t) \subseteq H^t$ be the set of private histories consistent with the action history portion of $\mathbb{h}^t$ in the mediated game. At any private history $h_i^t \in H(\mathbb{h}^t)$:

$$\begin{aligned}\mathbb{v}_i^t(\mathbb{s}_i,\bar{\mathbb{s}}_{-i}) &= (1-\delta)E_\lambda\left[u_i(a'_i,a_{-i}^t)|a_i^t,\mathbb{h}^t\right] + \delta E_\lambda\left[\sum_{r=t+1}^{\infty}\delta^{r-t-1}u_i(a^r)\middle|(a'_i,\mathbb{h}^t)\right]\\ &= (1-\delta)u_i(a'_i,s_{-i}(\alpha|h_i^t)) + \delta\kappa_i^{t+1}((h_i^t,a'_i)|(s'_i,s_{-i}))\\ &= v_i(s'_i,s_{-i})\end{aligned}$$

Thus, if $\mathbb{s}'_i$ is a profitable deviation from that recommended by the device λ_s in the mediated game, then there exists a profitable deviation in $G_\psi(\delta)$. This would contradict the premise that

s is a sequential equilibrium profile in $G_\psi(\delta)$. Since any strategy in G^{λ_s} is equivalent to some non-profitable deviation in $G_\psi(\delta)$, the correlation device λ_s and the obedient strategy $\bar{\mathfrak{s}}$ constitute a PCE of $G(\delta)$. □

4.5 MPCE Inclusion is Tight: Proof of Theorem I.4

($\subseteq$): Fix a game $G_\psi(\delta)$, and consider a sequential equilibrium strategy profile s with payoff v. Construct a PCE that induces the same outcome as s. Absent a pre-play signal, first period actions owe to independent mixtures, and so the PCE recommends an independent mixture in the first period. Next, by Lemma I.1, the continuation values prescribed by the PCE are in V^*. Thus, the payoff v is Nash enforced on V^*. □

($\supseteq$): We want to show that for every payoff w in $B_{NE}(V^*)$, there exists a monitoring structure ψ and a sequential equilibrium s of $G_\psi(\delta)$ with the same payoff w. Consider one such payoff and the pair $(\mu, \Bbbk)$ that Nash enforces it on V^*. Thus, for every action profile $a^j \in A$, $j = 1, \ldots, |A|$, there is a payoff $w^j \in V^*$ that is enforced on V^* by the admissible pair $(\mu^j, \Bbbk^j)$. Let the monitoring structure ψ provide perfect monitoring of actions, as well as a vector of private signals for each player. In particular, after the action profile a^j, each player privately observes his component of a draw from μ^j. So defined, consider the following strategy profile s in $G_\psi(\delta)$. "In the first period, mix according to μ. Following every subsequent history, choose the action corresponding to the most recently received message." Since the private messages are MPCE recommendations, s constitutes a Nash equilibrium. Then there exists a sequential equilibrium with the same path as s. So, there exists a private monitoring sequential equilibrium with the payoff w. □

4.6 Sufficiency

Assume that $\mathcal{P}$ is more strategically informative than $\mathcal{P}$. Let $G \in \Gamma_n$ be an arbitrary game and $v \in \Pi(G, \mathcal{Q})$ an arbitrary Bayes-Nash equilibrium payoff attained by the strategy profile f. Since $\mathcal{P}$ is a refinement of $\mathcal{Q}$, the strategy profile f is measurable with respect to $\mathcal{Q}$. Furthermore, since $\mathcal{P}_i$ is conditionally independent of $\mathcal{Q}_{-i}$, f is still a Bayes-Nash equilibrium.

4.7 Necessity

Suppose that $\mathcal{P}$ is more strategically informative than $\mathcal{Q}$. Since $\Pi(G, \mathcal{Q}) \subseteq \Pi(G, \mathcal{P})$ for every game $G \in \Gamma_n$, the inclusion in particular holds for all *decision games*: a game G^{D_i} in which players $-i$ earn a payoff of zero for every action profile, and the remaining player i — the decision maker — has non-null payoffs. Equilibrium strategies are by definition known to the decision maker, and so this game is equivalent to a standard choice under uncertainty problem, where the players $-i$ (with null payoffs) take the role of Nature. Let $\Pi_i(G^{D_i}, \mathcal{P})$ be set of payoffs to admissible decision rules in G^{D_i}. By Blackwell's Theorem, the hypothesis $\Pi_i(G^{D_i}, \mathcal{Q}) \subseteq \Pi_i(G^{D_i}, \mathcal{P})$ for all decision games is equivalent to $\mathcal{P}_i$ being sufficient for $\mathcal{Q}_{-i}$. Then by the Fisher factorization theorem, $\mathcal{P}_i$ is sufficient for $\mathcal{Q}_{-i}$ if and only if $\mathcal{P}_i$ it can be factored into two terms, one of which is conditionally independent of $\mathcal{Q}$.

Bibliography

Abreu, D., P. Milgrom, and D. Pearce (1991): "Information and Timing in Repeated Partnerships," *Econometrica*, 59(6), 1713–1733.

Abreu, D., D. Pearce, and E. Stacchetti (1986): "Optimal Cartel Equilibria With Imperfect Monitoring," *Journal of Economic Theory*, 39, 251–269.

——— (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.

Ahn, T., R. M. Isaac, and T. C. Salmon (2009): "Coming and going: Experiments on endogenous group sizes for excludable public goods," *Journal of Public Economics*, 93(1-2), 336 – 351.

Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

——— (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55(1), 1–18.

Baye, M. R., and H. C. Hoppe (2003): "The strategic equivalence of rent-seeking, innovation, and patent-race games," *Games and Economic Behavior*, 44(2), 217 – 226.

Bennathan, E., and A. A. Walters (1969): "Revenue Pooling and Cartels," *Oxford Economic Papers*, 21(2), pp. 161–176.

Bhaskar, V., and I. Obara (2002): "Belief-Based Equilibria in the Repeated Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory*, 102, 40–69.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24(2), 265–272.

Brekke, K. A., K. E. Hauge, J. T. Lind, and K. Nyborg (2009): "Playing with the Good Guys - A Public Good Game with Endogenous Group Formation," .

Brown, R. W. (1994): "Incentives and revenue sharing in college football: Spreading the wealth or giving away the game?," *Managerial and Decision Economics*, 15(5), 471–486.

Buckley, N., S. Mestelman, R. A. Muller, S. Schott, and J. Zhang (2009): "Shut Up and Fish: The Role of Communication when Output-Sharing is used to Manage a Common Pool Resource," *McMaster University Department of Economics Working Papers*, 2009-15(2009-15).

Camerer, C. F. (2003): *Behavioral Game Theory: Experiments in Strategic Interaction.* Princeton University Press.

Canes, M. E. (1974): "The Social Benefits of Restrictions on Team Quality," in *Government and the Sports Business*, ed. by R. G. Noll. Brookings Institution.

Charness, G. B., and C.-L. Yang (2008): "Endogenous Group Formation and Public Goods Provision: Exclusion, Exit, Mergers, and Redemption," *University of California eScholarship Repository.*

CHUNG, T.-Y. (1996): "Rent-seeking contest when the prize increases with aggregate efforts," *Public Choice*, 87, 55–66.

DASGUPTA, P. S., AND G. M. HEAL (1985): *Economic Theory and Exhaustible Resources*, no. 9780521297615 in Cambridge Books. Cambridge University Press.

D'ASPREMONT, C., AND A. JACQUEMIN (1988): "Cooperative and Noncooperative R & D in Duopoly with Spillovers," *The American Economic Review*, 78(5), pp. 1133–1137.

DEBREU, G. (1952): "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences of the United States of America*, 38(10), 886–893.

DIETL, H. M., E. FRANCK, AND M. LANG (2008): "Overinvestment in team sports leagues: A Contest Theory model," *Scottish Journal of Political Economy*, 55(3), 353–368.

ELY, J., J. HORNER, AND W. OLSZEWSKI (2005): "Belief-free Equilibria in Repeated Games," *Econometrica*, 73(2), 377–415.

ELY, J., AND J. VALIMAKI (2002): "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory*, 102(1), 84–105.

FALK, A., E. FEHR, AND U. FISCHBACHER (2005): "Driving Forces Behind Informal Sanctions," *Econometrica*, 73(6), 2017–2030.

FISCHBACHER, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics*, 10, 171–178.

FORGES, F. (1986): "An Approach to Communication Equilibrium," *Econometrica*, 54(6), 1375–1385.

FUCHS, W. (2007): "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review*, 97(4), 1432–1448.

FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5), 997–1039.

FUDENBERG, D., AND J. TIROLE (1991): *Game Theory.* Cambridge: MIT Press.

GOSSNER, O. (2000): "Comparison of Information Structures," *Games and Economic Behavior*, 30(1), 44 – 63.

GREEN, E. J., AND R. PORTER (1984): "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, 52(1), 87–100.

HEINTZELMAN, M. D., S. W. SALANT, AND S. SCHOTT (2009): "Putting free-riding to work: A Partnership Solution to the common-property problem," *Journal of Environmental Economics and Management*, 57(3), 309 – 320.

HERNANDEZ-LERMA, O. (1989): *Adaptive Markov Control Processes.* Springer-Verlag.

JUDD, K., S. YELTEKIN, AND J. CONKLIN (2003): "Computing Supergame Equilibria," *Econometrica*, 71(4), 1239–1254.

KAGI, W. (2001): "The Tragedy of the Commons Revisited: Sharing as a Means to Avoid Environmental Ruin," *Institute for Economy and the Environment IWOE Discussion Paper*, (91).

KAMIEN, M. I., E. MULLER, AND I. ZANG (1992): "Research Joint Ventures and R&D Cartels," *The American Economic Review*, 82(5), pp. 1293–1306.

KANDORI, M. (1992): "The Use of Information in Repeated Games with Imperfect Monitoring," *The Review of Economic Studies*, 59(3), 581–593.

KANDORI, M. (2002): "Introduction to Repeated Games with Private Monitoring," *Journal of Economic Theory*, 102, 1–15.

KANDORI, M. (2008): "Weakly Belief-Free Equilibria in Repeated Games with Private Monitoring," mimeo.

KANDORI, M., AND I. OBARA (2006): "Efficiency in Repeated Games Revisited: The Role of Private Strategies," *Econometrica*, 74(3), 499–519.

KONRAD, K. (2009): *Strategies and Dynamics in Contests*. Oxford University Press.

LEDYARD, J. O. (1995): "Public Goods: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by A. Roth, and J. Kagel. Princeton University Press.

LEHRER, E. (1992): "Correlated Equilibria in Two-Player Repeated Games with Nonobservable Actions," *Mathematics of Operations Research*, 17(1), 175–199.

MACLEOD, W. B. (2003): "Optimal Contracting with Subjective Evaluation," *The American Economic Review*, 93(1), pp. 216–240.

OSTROM, E. (1990): *Governing the commons: the evolution of institutions for collective action*. Cambridge University Press.

OSTROM, E., AND J. WALKER (1991): "Communication In a Commons: Cooperation Without External Enforcement," in *Laboratory Research in Political Economy*, ed. by T. R. Palfrey. University of Michigan Press.

OSTROM, E., J. WALKER, AND R. GARDNER (1992): "Covenants With and Without a Sword: Self-Governance is Possible," *The American Political Science Review*, 86(2), pp. 404–417.

——— (1994): *Rules, games, and common-pool resources*. University of Michigan Press,.

PAGE, T., L. PUTTERMAN, AND B. UNEL (2005): "Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency," *The Economic Journal*, 115(506), 1032–1053.

PICCIONE, M. (2002): "The Repeated Prisoners Dilemma with Imperfect Private Monitoring," *Journal of Economic Theory*, 102, 70–83.

PLATTEAU, J.-P., AND E. SEKI (2000): "Community Arrangements to Overcome Market Failures: Pooling Groups in Japanese Fisheries," in *Market, Community,and Economic Development*, ed. by M. Aoki, and Y. Hayami. Clarendon Press.

PORTER, R. H. (1983): "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *The Bell Journal of Economics*, 14(2), pp. 301–314.

POTEETE, A., M. JANSSEN, AND E. OSTROM (2010): *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press.

PROKOPOVYCH, P. (2006): "Perfect Correlated Equilibria in Repeated Games," University of Michigan Manuscript.

PUTTERMAN, L., J. T., AND K. KAMEI (2010): "Public Goods and Voting on Formal Sanction Schemes: An Experiment," *Department of Economics University of Copenhagen Working Paper*.

ROCKAFELLAR, R. T. (1996): *Convex Analysis*. Princeton University Press.

SAHLINS, M. (1972): *Stone Age Economics*. Aldine de Gruyter.

SALANT, S. W., AND G. SHAFFER (1999): "Unequal Treatment of Identical Agents in Cournot Equilibrium," *The American Economic Review*, 89(3), pp. 585–604.

Schott, S., N. Buckley, S. Mestelman, and R. Muller (2007): "Output sharing in partnerships as a common pool resource management instrument," *Environmental and Resource Economics*, 37, 697–711, 10.1007/s10640-006-9062-7.

Sekiguchi, T. (1997): "Efficiency in the Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory*, 76, 345–361.

Stahl, D. (1991): "The Graph of Prisoners' Dilemma Supergame Payoffs as a Function of the Discount Factor," *Games and Economic Behavior*, 3, 368–384.

Stigler, G. (1964): "A Theory of Oligopoly," *Journal of Political Economy*, 72(1), 44–61.

Tyran, J.-R., and L. P. Feld (2006): "Achieving Compliance when Legal Sanctions are Non-deterrent," *Scandinavian Journal of Economics*, 108(1), 135–156.

Walker, J. M., R. Gardner, A. Herr, and E. Ostrom (2000): "Collective Choice in the Commons: Experimental Results on Proposed Allocation Rules and Votes," *The Economic Journal*, 110(460), 212–234.