

*Unattainable Payoffs for Repeated Games of Private Monitoring**

Josh Cherry[†] Lones Smith[‡]

Department of Economics
University of Michigan

June 29, 2009

Abstract

We bound from the outside the set of sequential equilibrium payoffs in repeated games of private monitoring. The novelty in our approach is treating private histories as endogenous correlation devices. To do this, we develop a tractable new solution concept for standard repeated games with perfect monitoring: *Markov Perfect Correlated Equilibrium* generalizes the operator approach of Abreu, Pearce, and Stacchetti (1990) in a natural way to allow for correlated strategies. We show that *for any monitoring structure*, the set of sequential equilibrium payoffs of the repeated private monitoring game is always contained within the set of Markov Perfect Correlated Equilibrium payoffs of the analogous repeated game. Further, we explore two senses in which this bound is tight.

*We are grateful for helpful comments from Pavlo Prokopovych (whom Lones advised in his 2006 Michigan PhD) and seminar participants at the Western Economic Association Meetings, Midwest Theory Conference, and the University of Michigan.

[†]email: jscherry@umich.edu. Josh is thankful for the financial support NSF IGERT grants 0114368 and 0654014.

[‡]email: lones@umich.edu. Lones is grateful for the ongoing financial support of the NSF that has supported this work.

1 Introduction

A repeated game is a stylized model of a long-term relationship. The most often applied solution concepts for repeated games are Subgame Perfect Equilibrium (SPE) and its extension to imperfect monitoring, Perfect Public Equilibrium (PPE). In both cases, equilibrium strategies depend only on commonly observed histories. This yields a nice 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 thereby characterized equilibrium payoffs using methods inspired by dynamic programming.

Many interesting applications lack a public signal and thus these solution concepts lose traction. The earliest and most pointed example, Stigler’s (1964) “secret price cuts” model, studies a repeated Bertrand oligopoly model where firms secretly choose prices and only observe their own demand. Thus, the only sequential equilibria in public strategies are repetitions of stage game Nash equilibria. But this misses the potential richness of the dynamic structure — for in a non-trivial sequential equilibrium, computing a best response often requires a complicated probabilistic inference. Any useful recursive approach is therefore far from obvious.

Private monitoring in repeated games induces correlated private histories. In this paper, we interpret these as endogenous correlation devices. So motivated, we first develop a new solution concept for standard infinitely repeated games with perfect monitoring, and then explore its implications for repeated games of private monitoring. *Markov Perfect Correlated Equilibrium* (MPCE) imposes a correlated equilibrium at the start of every subgame, and has a recursive character like PPE. We characterize its payoff set by extending the operator methods of APS, and also develop an algorithm to compute it. We then shift to repeated games with 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 stage game. This helps us deduce the tightest bound on equilibrium payoffs that is independent of the monitoring structure.

Our analysis sheds light on payoffs when the folk theorem does not apply — such as when interaction is not very frequent, or more importantly, when information revelation about unobserved actions inherently cannot be accelerated. For as shown by Abreu, Milgrom, and Pearce (1991), the discrete time folk theorem logic unjustifiably yields more informative monitoring with more frequent play. This paper seeks to shed light on repeated games with fixed discount factors, as in APS.

In our attack on the problem, we explicitly account for possibilities occasioned by correlated private histories. On any sequential equilibrium path, the continuation play in any period is a correlated equilibrium, where the private histories are endogenous correlation devices. Though incentives are harder to provide when actions are unobserved, this induced correlation may facilitate better coordination (as in Aumann 1974, 1987), and substantially augment the sequential equilibrium payoff set. Our approach in fact admits arbitrary correlation each period.

We begin with a standard infinitely repeated game of observed actions, and embellish it with an extensive-form correlation device that can generate any history-dependent private messages every period. We assume that messages are made public after players act. However, 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 1 characterizes the MPCE payoff set — it is compact, convex, and nondecreasing in the discount factor. Also, this payoff set contains all subgame perfect payoffs. Theorem 2 then adds a tractable, recursive algorithm for computing it.

We next 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 3 shows that the MPCE payoff set serves as an upper bound for the sequential equilibrium payoffs. We thus identify the certainly unattainable sequential equilibrium payoffs for repeated games of private monitoring. Notably this bound holds for *all* monitoring structures.

Theorem 4 asserts that our payoff upper bound is nevertheless tight in the following two senses. First, in a private monitoring game with a doubly infinite time horizon, as introduced in Phelan and Skrzypacz (2008), every MPCE payoff is a sequential equilibrium payoff for some monitoring structure. But unlike MPCE, a standard repeated game of private monitoring does not allow pre-play signals. Instead, to understand a repeated game with an initial period, we modify the MPCE concept. We compute the Nash equilibrium payoffs of all auxiliary games using continuation payoffs drawn from the MPCE set. Put another way, this applies the APS operator to our MPCE payoff set. Any payoff in the resulting set can be supported as a sequential equilibrium in a repeated game with some private monitoring structure. Under either interpretation, we obtain the tightest possible bound that makes no reference to the monitoring structure.

In contrast to the literature, which has proceeded from the *inside* by finding computable subclasses of equilibria, in this paper, we bound the sequential equi-

librium payoff set from the *outside*. Existing work has also been motivated by the folk theorem, studying the case of very patient players. The first analyses of private monitoring found nearly efficient equilibria by dispensing with all but a simple summary of past play. Roughly speaking, these “belief-based” approaches focus on the probability that private messages are misleading. 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 has been identified by Piccione (2002) and Ely and Valimaki (2002), and greatly extended by Ely, Horner, and Olszewski (2005). While the belief-free approach can only identify a strict subset of all sequential equilibrium payoffs, it often secures a folk theorem.

The paper is organized as follows. We gently begin with a motivational example to which we later return. 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, returning to our example. We then formally describe infinitely repeated games with private monitoring, and compare payoffs for MPCE of standard repeated games and sequential equilibrium payoffs. Here, we establish our payoff upper bound and show that it can be tight. All proofs are in the Appendix.

2 Motivational Example

Consider a repeated Bertrand duopoly game with just two prices, high and low. We assume a strong temptation to cheat and a large penalty for being undercut (Figure 1). The firms share the discount factor $3/4$, and so are not patient enough to support the high price in a subgame perfect equilibrium. Stahl (1991) shows that even with public randomization, 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 $(7/2, 7/2)$.

	p_H	p_L
p_H	(4,4)	(-13,20)
p_L	(20,-13)	(0,0)

Figure 1: Stylized Bertrand Duopoly

Next, suppose that firms privately observe a payoff irrelevant signal of consumer taste. Before the firms choose prices, each privately sees the demand for its product

sold in the color green (g) or blue (b). The outcomes $\{(g, g), (g, b), (b, g)\}$ occur with probabilities $(1/2, 1/4, 1/4)$. To simplify matters, assume players can access a public randomization device that draws a number z from a uniform distribution on $[0, 1]$.

Consider the strategy profile: “In phase 1, charge a high price after observing green, and a low price after blue. If both firms charge the same price, then repeat phase 1. If firm i charges a low price, proceed to phase 2- i . In phase 2- i , firm i charges a high price, and firm $-i$ mixes so that firm i gets an expected payoff of 0. If both firms charge a high price, stay in phase 2- i . Otherwise, return to phase 1.”

When the repeated Bertrand duopoly is extended by the color information, 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 charge a high price, a firm will do so 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 firms expect the payoff $15/4$. In phase 2- i , firm i expects a payoff of 0 and firm $-i$ expects $15/2$. The payoff $(15/4, 15/4)$ Pareto dominates the highest symmetric subgame perfect equilibrium payoff $(7/2, 7/2)$ attainable without the information. Also, $(15/4, 15/4)$ 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.

3 A Mediated Repeated Game

We begin with a repeated game of perfect monitoring $G(\delta)$, played in periods $1, 2, \dots$, and payoffs discounted by the factor $0 < \delta < 1$. Each period, every player $i \in N = \{1, 2, \dots, n\}$ 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.¹ Payoffs given the action profile a are $u(a) = (u_1(a), \dots, u_n(a))$. Let α_i denote the mixed action

¹Throughout, 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.

for player i that chooses action $a_i \in A_i$ with chance $\alpha_i(a_i)$. Abusing notation, $u(\alpha) = (u_1(\alpha), \dots, u_n(\alpha))$ 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 the Revelation Principle, restricting messages to recommendations is without loss of generality.² Players commonly observe the null history $\mathfrak{h}^1 = \emptyset$ before play begins. A history $\mathfrak{h}^t = (a^1, \tilde{a}^1, \dots, a^{t-1}, \tilde{a}^{t-1})$ is a complete record of all past *outcomes* in periods $1, 2, \dots, t-1$, i.e. pairs of action and recommendation profiles. The history \mathfrak{h}^t is commonly observed by all players at the start of period t . Let \mathbb{H}^t be the set of all histories \mathfrak{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 $(\lambda^t : \mathbb{H}^t \rightarrow \Delta(A))_{t=1}^{\infty}$, and Λ is the space of all such functions.³ The interpretation is that after history \mathfrak{h}^t , the correlation device selects an action profile $\tilde{a} = (\tilde{a}_1, \dots, \tilde{a}_n) \in A$ according to the distribution $\lambda(\mathfrak{h}^t)$ 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 of the next history \mathfrak{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 s_i for player i is a sequence $(s_i^t)_{t=1}^{\infty}$, where $s_i^t : \mathbb{H}_i^t \times A_i \rightarrow \Delta(A_i)$ for every period $t = 1, 2, \dots$. So a strategy assigns a mixed action to every pair of history and recommendation. For any strategy profile $(s_1, \dots, s_n) = \mathfrak{s} \in \mathbb{S} = \prod_{i \in N} \mathbb{S}_i$, correlation device λ , and history \mathfrak{h}^t , the payoff for player i is the present value of future payoffs:

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

²In our finite model, the Revelation Principle holds since there cannot be issues with the measurable composition of functions.

³The 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.

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

$$v_i^t(s|h^t, \lambda) \geq v_i^t(\tilde{s}_i, s_{-i}|h^t, \lambda)$$

4 Markov Perfect Correlated Equilibrium

If $s \in \mathbb{S}$ is a sequential equilibrium strategy profile of $G^\lambda(\delta)$, then Prokopovych (2006) calls the pair (s, λ) 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 $\lambda : A^2 \rightarrow \Delta(A)$. Similarly, a strategy s is Markovian if it depends only on the most recent outcome and currently recommended action \tilde{a}_i , i.e. $s_i : A^2 \times A_i$. 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 (s, λ) 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 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 1 *Any PCE payoff is attainable in an MPCE.*

Because every MPCE is a PCE by definition, Lemma 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 $\mu = \lambda(h^t)$, 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* $\mathbb{k} : A^2 \rightarrow W$ 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 $(\tilde{a}_1, \dots, \tilde{a}_n) \in A$ is drawn from the distribution μ . Each player's type \tilde{a}_i

has the action set A_i , but the revised payoff function $E_\mu [(1 - \delta)u_i(a) + \delta k_i(a, \tilde{a}) | \tilde{a}_i]$ for the recommended action \tilde{a}_i .

If the distribution μ is a correlated equilibrium of $G_{\mathbb{k}}$, then the pair (μ, \mathbb{k}) is *admissible w.r.t. W* , where W is the co-domain of \mathbb{k} . (This is not to be confused with the range of \mathbb{k} , which is a subset of W .) In this case,

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

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 (μ, \mathbb{k}) is the (ex-ante) expected payoff $E_\mu [(1 - \delta)u(a) + \delta k(a)]$. Inversely, we write that the admissible pair (μ, \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

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

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 (μ^1, \mathbb{k}^1) supports w^1 and (μ^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 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 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 2 (Algorithm) *The MPCE payoff set is $V^* = \lim_{j \rightarrow \infty} W^j$, where the payoff set W^0 obeys $V^* \subseteq W^0$, and define $W^{j+1} = B(W^j)$ for $j = 1, 2, 3, \dots$*

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.

Let's return to the repeated Bertrand duopoly game of Section 2. In Figure 2, one can see that the MPCE payoff set is significantly larger than that of subgame perfect equilibrium. Appendix A.6 proves that the extreme feasible and individually rational payoffs $(132/17, 0)$ and $(0, 132/17)$ are also the highest asymmetric MPCE payoffs. So by convexity, the symmetric payoff $(66/17, 66/17)$ is also an MPCE, and in fact we argue that this is the highest symmetric MPCE payoff. This payoff is a convex combination of two extremal MPCE payoffs.




Figure 2: **Payoffs in the Repeated Bertrand Duopoly in Figure 1.** 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.

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 3, a repeated game is played in periods $1, 2, \dots$. Each period, every player $i \in N = \{1, 2, \dots, n\}$ 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 $m = (m_1, \dots, m_n)$ 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

$$u_i(a) = \sum_{m_i \in M_i} \psi_i(m_i|a) \pi_i(a_i, m_i) \quad (2)$$

We shall consider different monitoring structures ψ consistent with the same “expected stage game”. This requires that the payoffs $u(a) = (u_1(a), \dots, u_n(a))$ not depend on the monitoring structure. Since payoffs depend on ψ in (2), this exercise implies a corresponding change in the stochastic payoff structure π . Such a choice is possible provided (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.⁴ 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 $t = 1, 2, 3, \dots$. 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 $(a_i^1, \dots, a_i^{t-1})$ and past private messages $(m_i^1, \dots, m_i^{t-1})$, 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 $s_i^t : H_i^t \rightarrow \Delta(A_i)$ for every period $t = 1, 2, 3, \dots$. In other words, it maps every private into a mixed action. Let \mathcal{S} be the space of all such strategy profiles $s = (s_1, \dots, s_n)$.

⁴This 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.

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 $v_i : \mathcal{S} \rightarrow \mathbb{R}$ 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 $\tilde{s}_i : H_i \rightarrow \Delta(A_i)$ 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 mediated game $G_\psi(\delta)$.

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 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 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 $s \in \mathcal{S}$ that induce in $G^\lambda(\delta)$ the same outcome as does s in $G_\psi(\delta)$. After the history 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 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 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 in two senses.

Consider the repeated game $G_\psi^\infty(\delta)$ played in integer periods on the doubly infinite time horizon $\dots, -2, -1, 0, 1, 2, \dots$. Let V_ψ^∞ be the period 0 (say) sequential equilibrium payoff set in this new repeated game. The signal generated by the infinite history can provide no more than the information structure at the start of an MPCE. In fact, we show that the space of infinite histories can create any MPCE information structure.

Second, for a standard repeated game played in periods $1, 2, 3, \dots$, we can remove first period correlation from MPCE. An admissible pair (μ, \mathbb{k}) 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} :

$$B_{NE}(W) = \{v = E_\mu[(1 - \delta)\pi(a) + \delta\mathbb{k}(a, \tilde{a})] \mid (\mu, \mathbb{k}) \text{ is Nash admissible w.r.t. } W\}$$

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 4 (Tightness) (a) *Any MPCE payoff of $G(\delta)$ is attained in a sequential equilibrium of $G_\psi^\infty(\delta)$ for some monitoring structure ψ , and thus*

$$\bigcup_{\psi \in \Psi} V_\psi^\infty = V^*$$

(b) *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 2, Theorem 3 demonstrates that $(66/17, 66/17)$ is the highest symmetric sequential equilibrium in the infinitely repeated Bertrand duopoly 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 sequential equilibrium payoffs are bounded away from the Pareto frontier of the feasible and individually rational payoff set.

7 Conclusion

Understanding the equilibria of repeated games with private monitoring of actions has long been the next frontier in game theory. Yet finding sequential equilibria here has been hard, because recursive methods only capture subsets of them. In this paper, we have developed a new solution concept for repeated games, Markov Perfect Correlated Equilibrium, whose payoff set can be recursively computed. We have shown that its payoffs contain all equilibrium payoff sets for any repeated game of private monitoring, and that this bound is tight. Bounding payoffs from outside may prove fruitful in the future, and is currently a tractable tool for those in applied fields to use.

A Omitted Proofs

A.1 Any PCE payoff is an MPCE payoff: Proof of Lemma 1

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

We first argue that any pair admissible w.r.t. W is admissible to $\text{ext } W$.

Claim 1 (Bang-Bang) *For any pair (μ, \mathbb{k}) admissible w.r.t. the compact, convex set $W \subset \mathbb{R}^n$, there exists a pair $(\mu, \hat{\mathbb{k}})$ admissible w.r.t. $\text{ext } W$ with the same value.*

Proof of Claim 1: We closely parallel the proof of Theorem 3 in APS, accounting for correlation. For a bounded set $W \subset \mathbb{R}^n$, let $\mathcal{K}(W)$ be the set of all functions from A to W , and $\mathcal{K}(W, w) \subseteq \mathcal{K}(W)$ the set of continuation value functions that support w on W . The Banach-Alaoglu Theorem implies that the set $\mathcal{K}(W, w)$ is compact in the weak-* whenever W is compact. Further, $\mathcal{K}(W, w)$ is a convex set since a convex combination of admissible pairs is also an admissible pair. By the Krein-Milman Theorem, any $\hat{k} \in \mathcal{K}(W, w)$ can be written as a convex combination of extreme points of $\mathcal{K}(W, w)$. But by Proposition 6.1 of Aumann (1965), this is the same as the set of extreme points of $\mathcal{K}(\text{ext } W, w)$. So \hat{k} has the bang-bang property as it is a convex combination of extreme points of $\mathcal{K}(\text{ext } W, w)$. \square

Proof of Lemma 1: Let V_{PCE} be the set of PCE payoffs. Fix a PCE $\lambda \in \Lambda$ with payoff $w \in V_{PCE}$. 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. To do so, first rewrite the PCE λ as an admissible pair (μ, \mathbb{k}) . By Claim 1, the pair (μ, \mathbb{k}) is admissible w.r.t. at most $|A|^2$ points w_{jk} , $j, k = 1, \dots, |A|$, elements of $\text{ext } V_{PCE}$, each of which is enforced on the same set of extreme points. Let the pair $(\mu(a, \tilde{a}), \mathbb{k}(a, \tilde{a}))$ enforce $\mathbb{k}(a, \tilde{a})$. Then the Markovian device λ_M with $\lambda_M(a, \tilde{a}) = \mu(a, \tilde{a})$ is incentive compatible by construction. \square

A.2 Characterization of MPCE: Proof of Theorem 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 (μ, \mathbb{k}) that enforces w on W . Since W is self-generating, $\mathbb{k}(a, \tilde{a}) \in W$ for all outcomes (a, \tilde{a}) . Each payoff $\mathbb{k}(a, \tilde{a})$ 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 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, it suffices to show that V^* is self-generating. Consider an MPCE payoff $w \in V^*$. There exists a pair (μ, \mathbb{k}) such that $\mathbb{k}(a, \tilde{a}) \in V^*$ 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^* . Therefore, V^* is the largest fixed point of $B(\cdot)$. \square

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 (μ_j, \mathbb{k}_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 $\{\mu_l, \mathbb{k}_l\}$. The weak inequalities that define incentives are satisfied pointwise in the sequence $\{\mu_l, \mathbb{k}_l\}$, and so they are also satisfied in the limit (μ, \mathbb{k}) , which therefore enforces $b \in \mathbb{R}^n$. Therefore, $B(W)$ is closed since $b \in B(W)$. \square

Part (c) CONTAINS SUBGAME PERFECT EQUILIBRIUM PAYOFFS: Since the game has perfect monitoring of actions, players may ignore the correlation device, and instead play the subgame perfect equilibrium behavior after every history. \square

Part (d) NONDECREASING δ : We omit the proof as it is very similar to that of Theorem 6 of APS. \square

A.3 Algorithm: Proof of Theorem 2

We establish the validity of our algorithm by extending 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. 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 $k = 1, 2, \dots$. 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 $k = 1, 2, \dots$,

The sequence $\{W^k\}_{k=0}^\infty$ is bounded and monotone, and so converges (in the Hausdorff topology) to a point in the complete lattice \mathbb{W} . Let $W^\infty = \lim_{k \rightarrow \infty} W^k$. 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 1. \square

A.4 MPCE as an Upper Bound: Proof of Theorem 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

$$\kappa_i^t(h_i^t|s) = (1 - \delta)E \left[\sum_{r=t+1}^{\infty} \delta^{r-t-1} u_i(\beta_i^r) \middle| h_i^t, s \right] \quad (3)$$

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

$$(1 - \delta)u_i(\beta_i^t|h_i^t, s) + \delta\kappa_i^t(h_i^t|s) \geq (1 - \delta)u_i(\beta_i^t|h_i^t, \tilde{s}_i, s_{-i}) + \delta\kappa_i^t(h_i^t|\tilde{s}_i, s_{-i}) \quad (4)$$

for all players i , private histories h_i^t , and strategies \tilde{s}_i different from s_i in an information set.

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

Claim 2 (The Correlation Device) *For any strategy profile s of $G_\psi(\delta)$, there exists a correlation device $\lambda_s \in \Lambda$ and strategy \mathfrak{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 2: For any strategy profile $s \in \mathcal{S}$, let $\beta^t(a^t|(a^1, \dots, a^{t-1}), s)$ be the induced posterior probability of the action profile a^t in period t given the action history (a^1, \dots, 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,

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

For all action histories $\mathfrak{h}^t \in \mathfrak{H}^t$, define $\lambda_s(\mathfrak{h}^t) = \beta^t(\cdot|s, \mathfrak{h}^t)$. Then, after every action history, the recommendation distribution of λ_s coincides with the distribution of

actions in $G_\psi(\delta)$. Call \bar{s} the *obedient strategy* in $G^{\lambda_s}(\delta)$ — namely, where every player follows the recommendation of the correlation device λ after every history. Since the device λ_s recommends the same outcome as w , the obedient strategy \bar{s} in $G^{\lambda_s}(\delta)$ delivers the same outcome as s . \square

We must prove that obeying λ_s is a mutual best response for the players. This follows if $v_i^t(\bar{s}|\lambda_s) \geq v_i^t(s'_i, \bar{s}_{-i}|\lambda_s) \forall s'_i \in \mathcal{S}_i$. Our approach is to prove that for every deviation strategy $s'_i \in \mathcal{S}_i$ in the mediated game, there is a corresponding strategy $s'_i \in \mathcal{S}_i$ such that $v_i^t(s'_i, \bar{s}_{-i}|\lambda_s) = v_i^t(s'_i, s_{-i})$. Put differently, any deviation in the mediated game gives the same payoff as some private monitoring strategy; this cannot be a profitable deviation against the sequential equilibrium strategies s_{-i} . Thus, $v_i^t(\bar{s}|\lambda_s) = v_i^t(s) \geq v_i^t(s'_i, s_{-i}) = v_i^t(s'_i, \bar{s}|\lambda_s)$, as required.

Claim 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 3: By the one-shot deviation principle, the obedient strategy is a best reply to itself if and only if 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 the history $h^t \in H^t$ in which strategy s'_i instead plays the action a'_i in period t . Let $H(h^t) \subseteq H^t$ be the set of private histories consistent with the action history portion of h^t in the mediated game. At any private history $h_i^t \in H(h^t)$

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

So, if s'_i is a profitable deviation (from the recommendation of 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 $G_\psi(\delta)$. Thus, 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{s} constitute a PCE of $G(\delta)$. \square

A.5 MPCE Inclusion is Tight: Proof of Theorem 4

Proof of Part (a):

(\subseteq): Consider a sequential equilibrium strategy profile s with payoff $w \in V_\psi^\infty$, for some monitoring structure ψ . By the Revelation Principle, we can substitute the infinite history with the implied period one behavior prescribed by s . Let $\mu \in \Delta(A)$ be this (correlated) action distribution. We can replicate the outcome of s in a PCE by using μ in period one, and then appeal to Theorem 3 to capture continuation equilibria in PCE. So, $w \in V^*$. \square

(\supseteq): Consider a payoff $w \in V^*$, and its associated correlation device $\lambda \in \Lambda_M$. Let the monitoring structure ψ yield perfect monitoring of actions. Following the action profile a^j , $j = 1, \dots, |A|$, each player privately observes his component of a recommendation profile drawn from $\psi(a^j) = \lambda(a^j, a^j)$. Consider the equivalence class of strategy profiles, measurable w.r.t. to the Borel set of infinite histories in $G_\psi^\infty(\delta)$, that induce the distribution $\lambda(a^j, a^j)$ on action profiles $a^j \in A$ in period 0, i.e. the monitoring structure mimics the device assuming obedience. Select the strategy profile s from this class in which everyone chooses the action that coincides with his most recently received private message. Then s induces the same outcome path as the MPCE λ , and so is a Nash equilibrium. But only by a player's own deviation can he reach an information set off the equilibrium path. Thus, there exists a sequential equilibrium with the same path (see Proposition 3 of Sekiguchi, 1997). So there exists a private monitoring sequential equilibrium with the payoff w . \square

Proof of Part (b):

(\subseteq): This direction is proved with a very similar argument as the one presented proof of Theorem 3. Fix a game $G_\psi(\delta)$, and consider a sequential equilibrium strategy profile s with payoff v . First, construct a PCE that induces the same outcome as s . Absent a pre-play signal, first period actions are the result of independent mixtures, and so the PCE recommends an independent mixture in the first period. Next, by Lemma 1 the continuation values prescribed by the PCE are in V^* . Thus, the payoff v is Nash enforced on V^* . \square

(\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 (μ, \mathbb{k}) that Nash enforces it on V^* . Thus, for every action profile $a^j \in A$, $j = 1, \dots, |A|$, there is a payoff $w^j \in V^*$ that is enforced on V^* by the admissible pair (μ^j, \mathbb{k}^j) . Define the monitoring structure ψ as in the

proof of Part (a). 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.” The private messages are MPCE recommendations, so s constitutes a Nash equilibrium. Just as in the proof of Part (a), there exists a sequential equilibrium with the same path as s . So, there exists a private monitoring sequential equilibrium with the payoff w . \square

A.6 Derivation of the MPCE Frontier

First, we 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 $\{(p_H, p_H), (p_H, p_L), (p_L, p_H), (p_L, p_L)\}$, respectively, and $w_1, w_2 \in \mathbb{R}^2$ the continuation payoffs for players 1, 2. Given the stage game of Figure 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 $w_1, w_2 \leq 132/17$, 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 charge a high price:

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

Solving this linear 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 the payoff $(66/17, 66/17)$ is the highest symmetric MPCE payoff, and thus is the tight upper bound on the symmetric payoffs in any sequential equilibrium of $G_\psi(\delta)$.

References

- 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.
- 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.
- BHASKAR, V., AND I. OBARA (2002): “Belief-Based Equilibria in the Repeated Prisoners’ Dilemma with Private Monitoring,” *Journal of Economic Theory*, 102, 40–69.
- 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.
- FORGES, F. (1986): “An Approach to Communication Equilibrium,” *Econometrica*, 54(6), 1375–1385.
- FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62(5), 997–1039.
- 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.
- 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.
- PHELAN, C., AND A. SKRZYPACZ (2008): “Private Monitoring with Infinite Histories,” under revision.
- PICCIONE, M. (2002): “The Repeated Prisoners Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory*, 102, 70–83.
- PROKOPOVYCH, P. (2006): “Perfect Correlated Equilibria in Repeated Games,” University of Michigan Manuscript.
- 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.