A Folk Theorem with Mediated Communication

David Rahman*
University of Minnesota

PRELIMINARY AND INCOMPLETE

First Draft: March 31, 2011
This Version: March 2, 2012

Abstract

I study repeated games with mediated communication, with two main results: (i) a simple, intuitive algorithm for computing equilibrium payoffs as the discount factor tends to one, and as an application, (ii) a tight Folk Theorem with minimal detectability conditions. These results apply whether monitoring is public, private, one-sided, conditionally independent or dependent. Furthermore, the limit set of payoffs consists of all communication equilibria, be they public or private.

JEL Classification: D21, D23, D82.
Keywords: folk theorem, private strategies, mediated communication.

*I thank V. Bhaskar for useful conversations, Ichiro Obara for discussions on a related project, and especially Jerome Renault, Takuo Sugaya and Tristan Tomala for pointing out errors in a previous draft. I also thank the National Science Foundation for financial support through Grant No. SES 09-22253.

1 Introduction

I study repeated games with mediated communication, with two main results: (i) a tight Folk Theorem with minimal detectability conditions and (ii) a simple, intuitive algorithm for computing equilibrium payoffs as the discount factor $\delta \to 1$. These results apply whether monitoring is public, private, one-sided, conditionally independent or dependent. Furthermore, the limit set of payoffs consists of all communication equilibria, be they public or private.

The algorithm reduces the question of finding *all* communication equilibrium payoffs, including private equilibria, to an intuitive one-shot contracting problem. This weakens substantially the algorithm of Fudenberg and Levine (1994), which was limited to public equilibria. The Folk Theorem exploits this algorithm to derive weak detectability conditions for attaining feasible, individually rational payoffs as $\delta \to 1$.

To intuitively describe the algorithm, assume public monitoring for the time being and recall the algorithm of Fudenberg and Levine (1994). For each vector of welfare weights $\lambda \neq 0$ and given a mixed strategy profile $\mu \in \prod_i \; A_i$, define the *score* in the direction λ as follows: $k(\lambda) = \sup_\mu k(\lambda|\mu)$ and

$$k(\lambda|\mu) = \max_x \sum_{(i,a)} \lambda_i \mu(a)[u_i(a) - \sum_{s \in S} \Pr(s|a) x_i(s)] \text{ s.t. } \sum_{i=1}^n \lambda_i x_i(s) \geq 0 \quad \forall s,$$

$$\sum_{a_{-i}} \mu(a) \quad u_i(a, b_i) \leq \sum_{(a_{-i},s)} \mu(a) x_i(s) \quad \Pr(s|a, b_i, \rho_i) \quad \forall (i, a_i, b_i).$$

The problem is to maximize welfare subject to a self-generation "budget constraint" and incentive compatibility, with payoffs that include the variable x, which is proportional to continuation values in the repeated game. The set of public equilibrium payoffs as $\delta \to 1$ consists precisely of $\bigcap_\lambda h(\lambda)$, where $h(\lambda) = \{v \in \mathbb{R}^n : v \cdot \lambda \leq k(\lambda)\}$ is the half-space in the direction λ generated by the score $k(\lambda)$.

In recent papers, Rahman (2008) and Rahman and Obara (2010) suggested the use of recommendation-contingent payments, or continuation values, which in the context of repeated games delivers an algorithm based on the following score.

$$\max_x \sum_{(i,a)} \lambda_i \mu(a)[u_i(a) - \sum_{s \in S} \Pr(s|a) x_i(a,s)] \text{ s.t. } \sum_{i=1}^n \lambda_i x_i(a,s) \geq 0 \; \forall (a,s),$$

$$\sum_{a_{-i}} \mu(a) \quad u_i(a, b_i) \leq \sum_{(a_{-i},s)} \mu(a) x_i(a,s) \quad \Pr(s|a, b_i, \rho_i) \quad \forall (i, a_i, b_i).$$

The score above differs from the previous one in two respects: (i) continuation values may depend on the mediator's recommendations, and (ii) self-generation constraints are indexed by (a, s) and not just s. To see why the second difference is there, recall that, since equilibria are public, at the end of each period the profile of recommendations is announced publicly before the next round of recommendations. As a result, the self-generation constraints must be satisfied for every possible realization of both the public signal and the recommendation profile.

The new algorithm proposed in this paper uses the following score for $\lambda \geq 0$:[1]

$$\max_x \sum_{(i,a)} \lambda_i \mu(a)[u_i(a) - \sum_{s \in S} \Pr(s|a) x_i(a,s)] \text{ s.t. } \sum_{(i,a)} \lambda_i \mu(a) x_i(a,s) \geq 0 \;\forall s,$$

$$\sum_{a_{-i}} \mu(a) \quad u_i(a, b_i) \;\leq \sum_{(a_{-i},s)} \mu(a) x_i(a,s) \quad \Pr(s|a, b_i, \rho_i) \quad \forall (i, a_i, b_i).$$

The new algorithm differs from the old one in that the self-generation constraints are indexed by s and not (a, s). Intuitively, continuation values are averaged out over all possible recommendations, so that self-generation constraints only need hold in expectation with respect to recommendations. This weaker family of self-generation constraints completely exhausts players' ability to improve social outcomes with equilibria that are public only across t-period blocks, for all t. Intuitively, self-generation can be smeared over time for every fixed public observation s. A similar algorithm applies to the case of private monitoring, which is discussed later.

The value of this new score is manifold. First, it is simple: it collapses the set of attainable payoffs from "t-public" equilibria for all t into a one-period contracting problem. Secondly, it is intuitive: allowing for arbitrarily long t-period blocks can be thought of as spreading the self-generation "budget" over all possible realized action profiles. Thirdly, it is useful: it is easy to derive sufficient and almost necessary conditions for the Folk Theorem from it. Finally, it is general: it easily extends to games with private monitoring (Section 4), and perhaps more interestingly, encompasses *all* communication equilibrium payoffs. This last advantage is noteworthy because the problem of understanding private equilibria is mostly unexplored in the literature.

[1] The algorithm associated with $\lambda \ngeq 0$ is only slightly different, and derived in Section 4.

2 Examples

In this section I present some motivating examples that illustrate the paper's main results. I begin with the classic Prisoners' Dilemma, followed by a variant of it.

2.1 Prisoners' Dilemma with Public Monitoring

Let us begin with the Prisoner's Dilemma with public monitoring.

Example 1. There are two players, each of whom may cooperate, C, or defect, D. Stage game payo s are standard and given by the left bi-matrix below; there are two public signals (g or b) with probabilities on the right:

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

Payo s

	C	D
C	$\frac{2}{3}, \frac{1}{3}$	$\frac{1}{3}, \frac{2}{3}$
D	$\frac{1}{3}, \frac{2}{3}$	$0, 1$

Probabilities

The stage game is repeated in nitely often. At each stage, players take an action and subsequently observe a public signal according to the conditional probabilities above. Players earn average discounted payo s with discount factor $\delta \in [0, 1)$, as usual.

In this example, the set of symmetric perfect public equilibrium payo s is bounded away from the e cient pro le $(1, 1)$ uniformly in the discount factor δ. This result is standard, see Green and Porter (1984), Radner et al. (1986), and Fudenberg and Levine (1994), from which the example originates. One way to see this is with the usual observation of imperfect monitoring. Another way to see this is by noticing that pairwise identi ability fails in a binding way. Fudenberg and Levine (1994, p. 1004) provide a particularly simple proof.

One might think that the approach to private strategies by Kandori and Obara (2006) might be able to help with this example. However, it is not di cult to see that their approach also fails here, since they require that (i) the probability of \good news" decreases with the number of deviators, and (ii) the decrease is increasing. A more general approach is that derived by Rahman and Obara (2010, Example 1), which is a starting point for the present paper. I replicate their simple, intuitive solution below: players use a mediator to secretly take turns being a *secret principal*.

Consider the following correlated strategy: A mediator secretly recommends both players to play their part of (C, C) with probability $1 - 2\mu$, as well as both (C, D) and (D, C) each with probability $\mu > 0$. Continuation values, denoted by w, will be contingent not only on the public signal, but also on recommendations as follows.

$$
\begin{aligned}
w_1(CCs) &= w_2(CCs) = v, \\
w_1(CDs) + w_2(CDs) &= 2v, \\
w_1(CDs) = w_2(DCs), \quad &\text{and} \quad w_1(DCs) = w_2(CDs),
\end{aligned}
$$

where v is each player's equilibrium payoff (to be determined) and s is any signal realization, i.e., either g or b. It is easy to see using these inequalities that the equilibrium payoff of each player is given by

$$
\begin{aligned}
v &= (1-\delta)(1-\mu) + \delta[(1-2\mu)v + \mu(\tfrac{2}{3}w_1(CDb) + \tfrac{1}{3}w_1(CDg)) \\
&\quad + \mu(\tfrac{2}{3}w_1(DCb) + \tfrac{1}{3}w_1(DCg))] \\
&= 1 - \mu.
\end{aligned}
$$

Intuitively, players are first paid a constant amount v regardless of the signal, so there are no incentives. After (C, C), no \money" changes hands, so no incentives are created. After (C, D) is recommended (the same occurs with player roles reversed after (D, C)), we pick continuation values such that player 2 pays player 1 if the signal is g, and player 1 pays player 2 if the signal is b. Let $= w_1(CDg) - w_1(CDb) \geq 0$. Each player faces two incentive constraints: (i) playing D when asked to play C and (ii) vice versa. The first incentive constraint looks like

$$
(1-\delta)(1-\mu) \leq \delta\mu\tfrac{1}{3} \qquad \Rightarrow \qquad \geq \frac{3(1-\delta)(1-\mu)}{\delta\mu}.
$$

After applying the restrictions on w above, the second incentive constraint looks like

$$
-(1-\delta) \leq \delta\tfrac{1}{3} \qquad \Rightarrow \qquad \geq -\frac{3(1-\delta)}{\delta},
$$

so the second constraint is implied by the first. Therefore, by picking to satisfy the first constraint, we therefore found an intuitive dynamic, recursive arrangement that attains an equilibrium payoff of $1 - \mu$ for each player as $\delta \to 1$. Letting $\mu \to 0$, the Folk Theorem now follows. Therefore, a secret principal can solve the partnership problem of Radner et al. (1986) and Fudenberg and Levine (1994) in communication equilibrium, and eliminate the discontinuity in the equilibrium payoff correspondence with respect to the discount factor δ as $\delta \to 1$. The reason, as explained in Rahman and Obara (2010), is that every profitable deviation profile is attributable.

Now consider the following related example.

Example 2. The same as Example 1 except that the probability vector over signals conditional on (D, D) is now $(\frac{1}{2}, \frac{1}{2})$ instead of $(0, 1)$.

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

Payoffs

	C	D
C	$\frac{2}{3}, \frac{1}{3}$	$\frac{1}{3}, \frac{2}{3}$
D	$\frac{1}{3}, \frac{2}{3}$	$\frac{1}{2}, \frac{1}{2}$

Probabilities

In this example we cannot use the approach of Example 1, because there is a profitable, unattributable deviation profile. Indeed, if each player plays D when asked to play C and plays C with probability $\frac{1}{2}$ when asked to play D then conditional on any action profile each player changes the probability over signals the same way.

Notice, however, that player 1's deviation changes the probability over signals differently when player 2 plays C from when player 2 plays D. On the other hand, in Example 1, the change was always the same regardless of what player 2 played, but every deviation profile was attributable. This is a crucial observation that will deliver the Folk Theorem. If every disobedience is detectable, a detectable deviation profile is either attributable or for each player i there exist two action profiles for the others such that i's deviation changes the probability over signals differently. This permits a combination of the techniques by Abreu et al. (1990a) and Rahman (2008), where efficiency losses are eliminated asymptotically by delaying the arrival of effort recommendations.

2.2 Prisoners' Dilemma with Private Monitoring

Now consider the Prisoners' Dilemma with private monitoring.

C C

Example 3 (*One sided monitoring*). Just as in Example 1 above, there are two players, each of whom may cooperate, C, or defect, D. Stage game payo s are standard and given by the left bi-matrix below; signal probabilities are on the right:

	C	D
C	$1, 1$	$-1, 2$
D	$2, -1$	$0, 0$

Payo s

	C	D
C	$\frac{3}{4}, \frac{1}{4}$	$\frac{1}{2}, \frac{1}{2}$
D	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{2}$

Probabilities

The di erence here is that now the signal is only observed by player 1. Player 2 observes nothing (apart from his own actions and recommendations). Also notice that the probabilities are slightly di erent. Player 1 paying D leads to an uninformative signal, but playing C is informative of player 2's action.

This example is interesting for several reasons. Firstly, one sided monitoring is completely excluded by the literature on repeated games. Secondly, as will be seen from the general model, every deviation from the e cient outcome is detectable, so it can be approximated by equilibria of the repeated game. However, for the same reasons as for Example 1, T-private equilibria are needed to do so.

Example 4 (*Two reaso1an*

be approximated .0400..335 Td [52352 0 Td [(.)]TJ42 g 0 G2F17 11.9552 Tf 9.47 0 Td [(J

each of whom may cooperate, C, or defect, D

C

2.3 Efficiency and Observability

The next examples show the limits of detectability when trying to sustain e cient behavior. On the one hand, Radner, Myerson and Maskin showed that there may be a discontinuity between what is attainable as $\delta \to 1$ and when $\delta = 1$, in some sense. Example 1 suggests that this is not the case. Indeed, as long as every deviation is detectable, the example suggests that there is a sequence of (perhaps private) equilibria whose payo s converge to e ciency, which Radner (1986) shows is attainable when $\delta = 1$. On the other hand, Lehrer (1992) shows that, when $\delta = 1$, an e cient outcome is attainable essentially if and only if every pro table deviation from the e cient outcome is detectable by some not necessarily e cient action pro le.

The examples below show that this condition does not characterize attainability as $\delta \to 1$. In each of the examples below, the e cient outcome (D, C) cannot be approximated as $\delta \to 1$, yet every pro table deviation from that outcome is detectable by some action pro le, not necessarily the e cient action pro le.

Example 5 (*Detecting pro table deviations*). There are two players. Stage game payo s appear in the left bi-matrix below; public signal probabilities are on the right:

	C	D	E
C	$2, 2$	$0, 3$	$0, 2$
D	$3, 0$	$1, 0$	$1, 1$

Payo s

	C	D	E
C	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{2}$
D	$\frac{3}{4}, \frac{1}{4}$	$\frac{1}{4}, \frac{3}{4}$	$\frac{3}{4}, \frac{1}{4}$

Probabilities

In this example there exists no equilibrium such that the e cient outcome (D, C) is attained even virtually as $\delta \to 1$. To see this, rst notice that that the pro le (D, C) cannot be played with probability one. Otherwise, player 2 can pro tably deviate to play D instead of C without being detected. Furthermore, the pro le (D, C) cannot be played with positive probability. Otherwise, if player 1 plays C with positive probability then player 2 can play E which weakly dominates C and is also completely undetectable. On the other hand, for a deviation from (D, C) to be pro table, it must be the case that player 2 plays D with positive probability, but this is detectable when player 1 plays C. Therefore, although it is possible to sustain the payo pro le $(2, 2)$ in equilibrium when $\delta = 1$ in the sense of Lehrer (1992), it is impossible to do so as $\delta \to 1$.

For future reference, the same problem occurs even if we replace the probabilities given (C, D) with $(\frac{5}{6}, \frac{1}{6})$ instead. With these new probabilities, detection implies

attribution in the sense of Rahman and Obara (2010), because deviations from (C,C) shift probability in different directions. As will be argued later, attributability is important for understanding the set of payoffs attainable in a repeated game. The same problem occurs in the next example, although for a slightly different reason.

Example 6 (*Dominated detection*)**.** There are two players. Stage game payoffs appear in the left bi-matrix below; public signal probabilities are on the right:

	C	D	E
C	$2,2$	$0,3$	$0,2$
D	$3,0$	$1,0$	$1,1$
E	$2,0$	$0,0$	$0,0$

Payoffs

	C	D	E
C	$\frac{1}{2},\frac{1}{2}$	$\frac{1}{2},\frac{1}{2}$	$\frac{1}{2},\frac{1}{2}$
D	$\frac{3}{4},\frac{1}{4}$	$\frac{1}{4},\frac{3}{4}$	$\frac{3}{4},\frac{1}{4}$
E	$\frac{3}{4},\frac{1}{4}$	$\frac{1}{4},\frac{3}{4}$	$\frac{3}{4},\frac{1}{4}$

Probabilities

Again, although efficiency can be sustained when $\delta = 1$, it cannot be sustained as $\delta \to 1$. To see this, notice that although if player 1 plays E then the argument of the previous example fails, E is strictly dominated by C for player 1, and completely undetectable. Hence player 1 will not play E, so (D,C) is not attainable.

These two examples reflect all that can go wrong in trying to sustain efficiency in a repeated game due to failure of virtual enforceability, as I argue later when I develop the general model. However, one more thing could go wrong above and beyond virtual enforceability: not being able to enforce an action profile that changes probabilities differently for an unattributable deviation profile, as in the next example.

Example 7 (*Enforceability is insufficient*)**.** There are two players. Stage game payoffs appear in the left bi-matrix below; public signal probabilities are on the right:

	C	D	E
C	$1,1$	$-1,2$	$0,1$
D	$2,-1$	$0,0$	$-1,1$
E	$1,0$	$1,-1$	$\frac{1}{2},\frac{1}{2}$

Payoffs

	C	D	E
C	$\frac{2}{3},\frac{1}{3}$	$\frac{1}{3},\frac{2}{3}$	$\frac{2}{3},\frac{1}{3}$
D	$\frac{1}{3},\frac{2}{3}$	$0,1$	$\frac{1}{3},\frac{2}{3}$
E	$\frac{2}{3},\frac{1}{3}$	$\frac{1}{3},\frac{2}{3}$	$\frac{2}{3},\frac{1}{3}$

Probabilities

The payoff profile $(1,1)$ is not attainable, not even asymptotically, even though (C,C) is enforceable. If only (C,C) is played, defecting is unattributable, so everyone must be punished to discourage defection, which leads to inefficiency. On the other hand, the approach of Example 1 is not applicable here because if player 1 plays D with positive probability then player 2 can play E, which is weakly dominant and undetectable. Therefore, failure of attribution — not just detection — precludes the Folk Theorem in this example.

3 Model

Consider a repeated game with private monitoring. The stage game consists of a finite set $I = \{1, \ldots, n\}$ of players, a finite set A_i of actions for each player $i \in I$, where $A = \prod_i A_i$, utility profile $u : I \times A \to \mathbb{R}$, where $u_i(a)$ is the utility to player i from action profile a, and a finite set of signals S_i. Let $S = \prod_i S_i$ and $\Pr(s|a)$ be the probability that the profile s of signals is observed if the profile a of actions was played. A reporting strategy is any map $\rho_i : S_i \to S_i$, where $\rho_i(s_i)$ is what i reports to the mediator after observing s_i. Let

$$\Pr(s|a_{-i}, b_i, \rho_i) = \sum_{t_i \in \rho_i^{-1}(s_i)} \Pr(s_{-i}, t_i|a_{-i}, b_i)$$

be the probability of reported signals if everyone is honest and obediently playing a_{-i} except for i, who instead plays b_i and reports according to ρ_i.

Players have a common discount factor δ, where $\delta \in [0, 1)$. Given a sequence of action profiles $a^\infty = (a_1, a_2, \ldots)$, the utility to player i is given by

$$U_i(a^\infty) = [(1-\delta)/\delta] \sum_{t=1}^{\infty} \delta^t u_i(a_t).$$

Let $U = \text{conv}\{u(a) \in \mathbb{R}^n : a \in A\}$ be the convex hull of possible payoff vectors in the stage game. I make the following two relatively standard assumptions.

Assumption 1 (Full support). $\Pr(s|a) > 0$ for all (a, s).

This assumption means that every signal is possible, regardless of what players play. It is largely made for simplicity. This way "the entire tree lights up," in other words, nothing is off the path of play. I relax this assumption, but only a little bit, in Section ??, and make connections to some of the literature.

Let $\underline{u} = (\underline{u}_1, \ldots, \underline{u}_n)$ be the vector of correlated minmax values, where

$$\underline{u}_i = \min_{\mu \in \Delta(A)} \max_{\beta_i : A_i \to A_i} \sum_{a \in A} \mu(a) u_i(\beta_i(a_i), a_{-i}) \quad \forall i.$$

Let $U_+ = \{u \in U : u \geq \underline{u}\}$ be the set of *feasible, individually rational payoffs*.

Assumption 2 (Full dimensionality). $\dim U_+ = n$.

This standard assumption allows for players to be punished and rewarded independently. It allows for a version of the algorithm by Fudenberg and Levine (1994) to characterize equilibrium payoffs. It is easy to completely relax Assumption 2 following Fudenberg et al. (2007) and let U_+ have arbitrary dimension — see Section ??.

3.1 An Algorithm for Games with Public Monitoring

For now, let me make two simplifying assumptions: (i) public monitoring, i.e., players report signals truthfully and publicly, and (ii) constant support, i.e., supp $\mu = \prod_i B_i$ for some $B_i \subset A_i$.[2] Consider the problem of minimizing expected payments ξ that enforce a given correlated strategy μ subject to ex post budget-feasibility.

$$c_1(\mu|\lambda) = \min_{\xi} \sum_{(i,a,s)} \lambda_i \xi_i(a,s) \text{ s.t. } \sum_{i=1}^{n} \lambda_i \xi_i(a,s) \geq 0 \qquad \forall (a,s),$$

$$\sum_{a_{-i}} \mu(a) \quad u_i(a,b_i) \leq \sum_{(a_{-i},s)} \xi_i(a,s) \quad \mathrm{Lr}(s|a,b_i) \qquad \forall (i,a_i,b_i),$$

where a ranges over supp μ and $\mathrm{Lr}(s|a,b_i) = [\Pr(s|a_{-i},b_i)/\Pr(s|a)] - 1$. Intuitively, think of $\xi_i(a,s) = x_i(a,s)\mu(a)\Pr(s|a)$, for some monetary payments $x_i(a,s)$. The contractual problem consists of minimizing the weighted sum of expected payments subject to the following two constraints: (i) the weighted sum of payments state-by-state is non-negative (called \self-generation"), and (ii) the payments enforce μ.

This problem helps to de ne the *score* of a repeated game in every direction λ. Let $c_1^*(\lambda) = \sup_\mu \lambda \cdot u(\mu) - c_1(\mu|\lambda)$, and $\partial c_1^*(0) = \{v \in \mathbb{R}^n : \lambda \cdot v \leq c_1^*(\lambda) \ \forall \lambda \in \mathbb{R}^n\}$, i.e., the subdi erential of c_1^* at 0. Fudenberg and Levine (1994) showed that the set $\partial c_1^*(0)$ coincides with the limit set of public equilibrium payo s as $\delta \to 1$ when $\dim \partial c_1^*(0) = n$. This observation delivers their well-known algorithm for computing the limit set of public equilibrium payo s in terms of the \conjugate" function c_1^*.

Rather than view the repeated game as a repetition of the one-period stage game, one may view it as a repetition of a T-period stage game, for any T-period block, and apply the previous algorithm. Doing so has interesting implications for the self-generation constraint. De ne the following T-period contractual problem.

$$c_T(\mu|\lambda) = \min_{\xi,\Delta v} \sum_{(i,a^T,s^T)} \lambda_i \xi_i(a^T,s^T) \text{ s.t. } \sum_{i=1}^{n} \lambda_i \xi_i(a^T,s^T) \geq 0 \quad \forall (a^T,s^T),$$

$$v_{i\tau}(a_i^\tau, b_i^\tau, s^{\tau-1}) = \tfrac{1}{T} \sum_{a_{-i}^\tau} \Pr(s^{\tau-1}|a_{-i}^{\tau-1}, b_i^{\tau-1})\mu(a^\tau) \quad u_i(a_{-i\tau}, b_{i\tau})$$

$$- \sum_{(a^T,s^T)\geq(a_i^\tau,s^{\tau-1})} \xi_i(a^T,s^T) \quad \mathrm{Lr}(s_\tau|a_\tau,b_{i\tau})\mathrm{Lr}(s^{\tau-1}|a^{\tau-1},b_i^\tau) \quad \forall (i,\tau,a_i^\tau,b_i^\tau,s^{\tau-1}),$$

$$\sum_{(\tau,s^{\tau-1},a_i^\tau)} v_{i\tau}(a_i^\tau, \beta_i^\tau(a_i^\tau,s^{\tau-1}),s^{\tau-1}) \leq 0 \quad \forall (i,\beta_i^T),$$

[2]Private monitoring and non-constant support are simple extensions, developed in Section 4.

where τ ranges from 1 to T, $\beta_i^\tau = (\beta_{i1}, \ldots, \beta_{i\tau})$, and $\beta_{i\tau}(a_i^\tau, s^{\tau-1})$ is interpreted as the action taken after being recommended a_i^τ and having observed $s^{\tau-1}$.

The main result of the paper is that, for \irrevocable" games, defined below, $c_T(\mu|\lambda)$ converges to the value of the following *asymptotic contractual problem* (when $\lambda \geq 0$).

$$c_\infty(\mu|\lambda) \;=\; \min_{\xi} \sum_{(i,a,s)} \lambda_i \xi_i(a,s) \;\text{ s.t. }\; \sum_{(i,a)} \lambda_i \xi_i(a,s) \geq 0 \qquad \forall s,$$

$$\sum_{a_{-i}} \mu(a) \quad u_i(a,b_i) \leq \sum_{(a_{-i},s)} \xi_i(a,s) \quad \mathrm{Lr}(s|a,b_i) \qquad \forall (i,a_i,b_i).$$

The key difference between c_1 and c_∞ is that the self-generation constraint integrates out the action profiles in c_∞, whereas it does not in c_1. Therefore, at least in principle, the constraints in c_∞ are much weaker than in c_1. The problem c_∞ may be thought of as an extension of the insight by Abreu et al. (1990a) that delaying the arrival of news can economize on the cost of providing incentives. The main driver of this result is that the news of just which actions were recommended to players can be delayed in order to economize on incentive costs. To see this, it is helpful to look at the dual asymptotic problem.

$$c_\infty(\mu|\lambda) = \max_{\alpha,\eta \geq 0} \sum_{(i,a,b_i)} \mu(a) \quad u_i(a,b_i)\alpha_i(a_i,b_i) \text{ s.t. } \quad \mathrm{Lr}(s|a,\alpha_i) = \lambda_i(1-\eta(s)) \;\forall (i,a,s).$$

Heuristically, this dual maximizes the sum of players' deviation gains with deviation profiles that are (i) unattributable,[3] and (ii) generate the same change in likelihood ratios, for every recommended action profile (call them supp μ*-hampering* deviation profiles). As a result, delaying the arrival of recommendations to players does not economize on incentives. It turns out that these are in fact the only deviations whose gains cannot be mitigated by delaying arrival of news about recommendations.

The function c_∞ deals with non-negative welfare weights. To deal with the possibility of some welfare weights being negative, we must define another function, which resonates with the notion of \joint rationality" by Renault and Tomala (2004). Define

$$c_0(\mu|\lambda) \;=\; \max_{\sigma \geq 0} \; -\sum_{i:\lambda_i<0} \lambda_i \quad u_i(\mu,\sigma_i) \;\text{ s.t. }\; \sum_{b_i} \sigma_i(b_i|a_i) = 1 \qquad \forall (i,a_i,b_i),$$

$$\mathrm{Lr}(s|a,\sigma_i) = \lambda_i(1-\eta(a,s)) \qquad \forall (i,a,s),$$

except that $c_0(\mu|\lambda) = +\infty$ if μ is not enforceable.

[3] "Unattributable" means that $\Pr(s|a,\sigma_i)$ does not depend on i (see Rahman and Obara, 2010).

To help digest c_0 for a moment, notice that if $\lambda = -\mathbf{1}_i$ then $-c_0^*(\lambda)$ equals player i's minmax payoff: $c_0^*(\lambda) = -\min_\mu u_i(\mu) + \max_{\sigma_i} \quad u_i(\mu, \sigma_i) = -\min_\mu \max_{\sigma_i} u_i(\mu, \sigma_i)$.

We are almost ready to present the main result of the paper, but first I need one more definition. I will discuss this definition at greater length a few pages from now. Let $\mathscr{C}_i(a_i, s) = \text{cone}\{\text{Lr}(s|a_i, \cdot, b_i) : b_i \in \text{supp } \sigma_i(b_i|a_i)\}$ and define its *polar cone* as $\mathscr{C}_i(a_i, s)^\circ = \{\xi_i(a_i, \cdot, s) : \sum_{a_{-i}} \xi_i(a, s)\text{Lr}(s|a, b_i) \geq 0 \ \forall b_i \in \text{supp } \sigma_i(b_i|a_i)\}$.

Definition 1. A deviation profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is called supp μ*-revocable* if it is supp μ-unattributable and $\mathscr{C}_i(a_i, s)^\circ \cap (-\mathbf{1})^\circ \neq \{\mathbf{0}\}$, where $(-\mathbf{1})^\circ = \{x : \sum_k x_k \leq 0\}$. Otherwise, σ is called supp μ-irrevocable. A stage game is called supp μ*-irrevocable* if every supp μ-hampering deviation profile is supp μ-irrevocable. I will sometimes simply say "irrevocable" when the prefix is understood.

Although revocability looks like an obscure technical assumption at this moment, there is a simple way to interpret it: it is impossible to dynamically "undo" the statistical vulnerability of a deviation. It is easy to construct sufficient conditions for irrevocability, such as the following three: (i) $\mathbf{1} \notin \text{cone}\{\text{Lr}(s|a_i, \cdot, b_i) : b_i \neq a_i\}$ for some (i, a_i, s), (ii) if σ is supp μ-unattributable then $|\{b_i \neq a_i : \sigma_i(b_i|a_i) > 0\}| = 1$ for some (i, a_i), and (iii) the stage game is B-scrutable for every doubleton $B \subset A$, which follows if for every pair $a, \hat{a} \in A$ there is a triple (i, a_i, s) such that for all b_i, either $\text{Lr}(s|a, b_i) \leq \text{Lr}(s|a_i, \hat{a}_{-i}, b_i)$ or $\text{Lr}(s|a, b_i) \geq \text{Lr}(s|a_i, \hat{a}_{-i}, b_i)$. Of course, this last condition implies B-revocability for every subset $B \subset A$, since revocability for singleton sets follows immediately. Although I view it as mainly a technical condition, it can sometimes be substantially restrictive, though, as I suggest later.

I will discuss revocability and its interpretation a few pages from now, but first, let me state and begin to prove the paper's main result, establishing convergence of the sequence of values of the t-period problems to the combined value of each of the two the two asymptotic problems described above, c_∞ and c_0.

Given any $t \in \mathbb{N} \cup \{0, \infty\}$, let $c_t^*(\lambda) = \sup_\mu \lambda \cdot u(\mu) - c_t(\mu|\lambda)$ be the *conjugate function*, or *score*, of c_t in the direction λ, and write $\partial c_t^*(0) = \{v \in \mathbb{R}^n : \lambda \cdot v \leq c_t^*(\lambda) \ \forall \lambda\}$ for the subdifferential of c_t^* at 0. For every game and $t \in \mathbb{N}$, $\partial c_t^*(0)$ is precisely the set of limits of all t-public communication equilibrium payoffs as $\delta \to 1$. The main result of this paper states that for every irrevocable game, the union of limit sets of t-public equilibrium payoffs for all $t \in \mathbb{N}$ can be understood succinctly and intuitively.

Theorem 1 (Algorithm). *For every irrevocable game,*

$$\bigcup_{t\in\mathbb{N}} \partial c_t^*(0) = \partial c_\infty^*(0) \cap \partial c_0^*(0).$$[4]

Remarkably, Theorem 1 shows that the set of all limits of equilibrium payoffs can be described by a simple, intuitive algorithm. The left-hand side corresponds to limiting equilibrium payoffs as $\delta \to 1$ from all possible t-public equilibria, for all t simultaneously. The right-hand side has two parts. The first part, $\partial c_\infty^*(0)$, corresponds to the set of outcomes that are enforceable but for which delaying the arrival of information does not add value to players. The second part, $\partial c_0^*(0)$, roughly corresponds to, but is more general than, Renault and Tomala's set of "jointly rational" payoffs, as I argue at some length later.[5] Describing how these objects emerge in the limiting equilibrium payoff set will become easier as we go over the proof.

3.2 Proof for Theorem 1

The proof proceeds in four steps. I begin by assuming that $\lambda \geq 0$. In Step 1, I show that $c_t \geq c_\infty$ for all $t \in \mathbb{N}$. Then, in Step 2, I show that revocability is sufficient and almost necessary for $c_1 > c_2$ whenever $c_1 > c_\infty$. In this step I discuss the notion of revocability as well as its interpretation in some length. In Step 3, I show that $c_t \to c_\infty$ as $t \to \infty$. Finally, in Step 4, I generalize these arguments to include $\lambda \ngeq 0$.

First, assume that $\lambda = (1, \ldots, 1)$ is the summer vector, and consider the 1-period dual problem below. It reveals the cost-efficient way of providing incentives: with a maximum likelihood ratio test.

$$\begin{aligned} c_1(\mu) &= \max_{\alpha,\eta\geq 0} \sum_{i=1}^{n} u_i(\mu,\alpha_i) \quad \text{s.t.} \quad \mathrm{Lr}(s|a,\alpha_i) = 1 - \eta(a,s) \quad \forall(i,a,s) \\ &= \max_{\sigma} \min_{(a,s)} \sum_{i=1}^{n} \frac{u_i(\mu,\sigma_i)}{\mathrm{Lr}(s|a,\sigma_i)} \quad \text{s.t.} \quad \mathrm{Lr}(s|a,\sigma_i) = \mathrm{Lr}(s|a,\sigma_j), \end{aligned}$$

where $\sigma = (\sigma_1, \ldots, \sigma_n)$ satisfies $\sigma_i \geq 0$ and $\sum_{(b_i,\rho_i)} \sigma_i(b_i,\rho_i|a_i) = 1$ for all (i, a_i). Intuitively, in the dual first players pick an unattributable deviation profile to maximize

[4] With *private monitoring* the same theorem holds except that $c_\infty(\mu|\lambda)$ is replaced with the indicator function of enforceability, i.e., 0 if μ is enforceable and $+\infty$ otherwise. More on this later.

[5] Intuitively, the difference is comparable to that between virtual enforceability (Rahman, 2008) and sustainability (Lehrer, 1992) after adding budget constraints.

the ratio of deviation gains to changes in likelihood ratio, and then the principal picks a profile $(\hat{a}, \hat{s})$ to minimize this ratio. This results in the lowest expected sum of payments to players subject to the sum of ex post payments being non-negative.

To obtain the equation above, notice first that without loss, $\sum_{b_i} \alpha_i(b_i|a_i) = \alpha_{i0}$ independently of a_i. If not, increase $\alpha_i(a_i, a_i)$ for every a_i with $\sum_{b_i} \alpha_i(b_i|a_i) < \alpha_{i0}$ until the left-hand side reaches α_{i0}. Clearly, this change does not affect the dual. Now write $\Delta u_i(\mu, \alpha_i) = \Delta u_i(\mu, \sigma_i)\alpha_0$ and $\Delta \mathrm{Lr}(s|a, \alpha_i) = \Delta \mathrm{Lr}(s|a, \sigma_i)\alpha_0$. A feasible dual solution α with $\sum_i \Delta u_i(\mu, \alpha_i) > 0$ and $\Delta \mathrm{Lr}(s|a, \alpha_i) < 1$ for all (a, s) cannot solve the dual, since otherwise it would be possible to raise α_0, increasing the value of the dual objective without violating any constraints. Therefore, some dual constraint — say at $(\hat{a}, \hat{s})$ — must bind at an optimum. In other words, $\Delta \mathrm{Lr}(\hat{s}|\hat{a}, \alpha_i) = \Delta \mathrm{Lr}(\hat{s}|\hat{a}, \sigma_i)\alpha_0 = 1$, so $\alpha_0 = 1/\Delta \mathrm{Lr}(\hat{s}|\hat{a}, \sigma_i)$. Plugging this into the objective yields $\sum_i \Delta u_i(\mu, \sigma_i)/\Delta \mathrm{Lr}(\hat{s}|\hat{a}, \sigma_i)$, and the equation follows.[6]

Next, I derive and analyze the dual for $t = 2$, and use it to begin to prove Theorem 1. The 2-period dual contracting problem is given by the following program.

$$
\begin{aligned}
c_2(\mu) = \max_{\alpha, \eta \geq 0} \; & \tfrac{1}{2} \sum_{(i,a,b_i)} \mu(a_1) \Delta u_i(a_1, b_{i1}) \alpha_{i1}(a_{i1}, b_{i1}) \\
& + \tfrac{1}{2} \sum_{(i,t,a^t,b_i^t,s^{t-1})} \mu(a_1) \Pr(s_1|a_{-i1}, b_{i1}) \mu(a_2) \Delta u_i(a_{-i2}, b_{i2}) \alpha_{i2}(a_i^2, b_i^2, s_1) \text{ s.t.} \\
\Delta \mathrm{Lr}(s_1|a_1, \alpha_{i1}) + \sum_{b_i^2} & \Delta \mathrm{Lr}(s_2|a_2, b_{i2}) \mathrm{Lr}(s_1|a_1, b_{i1}) \alpha_{i2}(a_i^2, b_i^2, s_1) = 1 - \eta(a^2, s^2), \\
& \sum_{b_{i2}} \alpha_{i2}(a_i^2, b_i^2, s_1) = \alpha_{i1}(a_{i1}, b_{i1}), \\
& \sum_{b_{i1}} \alpha_{i1}(a_{i1}, b_{i1}) = \alpha_{i0}.
\end{aligned}
$$

This problem chooses a dynamic deviation profile to maximize the sum of deviation gains subject to being unattributable and a similar bound on likelihood ratios as in c_1 above. For the first step, notice first that by revealed preference, clearly $c_\infty \leq c_1$, since any feasible payment scheme for the 1-period problem is feasible for the asymptotic problem. It is also clear by revealed preference that $c_2 \leq c_1$: any feasible payment scheme ξ^1 for the 1-period problem can be halved and applied to each stage of the 2-period problem, i.e., the payment scheme $\xi^2(a_1, a_2, s_1, s_2) = \frac{1}{2}\xi^1(a_1, s_1) + \frac{1}{2}\xi^1(a_2, s_2)$ yields the same value as the 1-period problem. To see that in addition $c_\infty \leq c_2$, assume that $0 < c_1(\mu) < \infty$. (The other cases are trivial.)

[6] Let us agree that zero divided by zero equals zero.

Let (α, η) be any optimal solution of the dual asymptotic problem. Now, to see that $c_\infty \leq c_2$, I will use (α, η) to construct a feasible solution α_2 of the 2-period problem whose value equals $c_\infty(\mu)$. Let σ be defined by $\sigma_i(b_i|a_i) = \alpha_i(a_i, b_i)/\sum_{\hat{b}_i} \alpha_i(a_i, \hat{b}_i)$. Without loss, $\eta(\overline{s}) = 0$ for some $\overline{s}$, just as in the 1-period dual problem. Define α_{i2} as follows: $\alpha_{i2}(a_i^2, b_i^2, s_1) = \alpha_{i0}\sigma_{i2}(b_{i2}|a_{i2}, s_1)\sigma_{i1}(b_{i1}|a_{i1})$, where $\alpha_{i0} = \Delta\mathrm{Lr}(\overline{s}|a_1, \sigma_{i1})$,[7]

$$\sigma_{i\tau}(b_{i\tau}|a_{i\tau}, s_{\tau-1}) = \gamma_{i\tau-1}(s_{\tau-1})\sigma_i(b_{i\tau}|a_{i\tau}) + (1 - \gamma_{i\tau-1}(s_{\tau-1}))[a_{i\tau}](b_{i\tau})$$

for $\tau = 1$ and 2, and $[a_{i\tau}]$ is Dirac measure and γ_{i0} and $\gamma_{i1}(s_1)$ satisfy

$$\gamma_{i1}(s_1) = \gamma_{i0}\frac{\Delta\mathrm{Lr}(\overline{s}|a_1, \sigma_i) - \Delta\mathrm{Lr}(s_1|a_1, \sigma_i)}{\mathrm{Lr}(\overline{s}|a_1, \sigma_i)\mathrm{Lr}(s_1|a_1, \sigma_{i1})}.$$

Clearly, $\gamma \geq 0$ and $\gamma_{i1}(s_1) = 0$ if $s_1 = \overline{s}$, or if $\gamma_i^0 = 0$. Let $\underline{s} \in \arg\max_s R(s)$, where

$$R(s) = \frac{\Delta\mathrm{Lr}(\overline{s}|a, \sigma_i) - \Delta\mathrm{Lr}(s|a, \sigma_i)}{\mathrm{Lr}(\overline{s}|a, \sigma_i)\mathrm{Lr}(s|a, \sigma_i)}.$$

It is not difficult to see that $R(\underline{s}) > 1$. Indeed, clearly $\underline{s} \in \arg\min_s \mathrm{Lr}(s|a, \sigma_i)$, since the numerator of $R(s)$ is decreasing and the denominator increasing with $\mathrm{Lr}(s|a, \sigma_i)$. Since σ_i is detectable, $\mathrm{Lr}(\underline{s}|a, \sigma_i) < 1$ and $0 < R(\underline{s}) < \infty$. Now, if $R(\underline{s}) \leq 1$ then

$$\begin{aligned}
\Delta\mathrm{Lr}(\overline{s}|a, \sigma_i) - \Delta\mathrm{Lr}(\underline{s}|a, \sigma_i) &\leq \Delta\mathrm{Lr}(\overline{s}|a, \sigma_i)\mathrm{Lr}(\underline{s}|a, \sigma_i) \\
\Leftrightarrow \mathrm{Lr}(\overline{s}|a, \sigma_i) - \mathrm{Lr}(\underline{s}|a, \sigma_i) &\leq \mathrm{Lr}(\overline{s}|a, \sigma_i)\mathrm{Lr}(\underline{s}|a, \sigma_i) - \mathrm{Lr}(\underline{s}|a, \sigma_i) \\
\Leftrightarrow \mathrm{Lr}(\overline{s}|a, \sigma_i) \leq \mathrm{Lr}(\overline{s}|a, \sigma_i)\mathrm{Lr}(\underline{s}|a, \sigma_i) \quad &\Leftrightarrow \quad 1 \leq \mathrm{Lr}(\underline{s}|a, \sigma_i),
\end{aligned}$$

a contradiction, so $R(\underline{s}) > 1$. Given γ_{i0}, all other $\gamma_{i1}(s_1)$ are implied by the defining formula above, but γ_{i0} remains undetermined. Choose γ_{i0} so that $\gamma_{i1}(\underline{s}) = 1$. Clearly this is possible, since (i) $\gamma_{i0} = 0$ implies that $\gamma_{i1}(\underline{s}) = 0$, (ii) $\gamma_{i0} = 1$ implies that $\gamma_{i1}(\underline{s}) > 1$, and (iii) $\gamma_{i1}(\underline{s})$ is continuous in γ_{i0}. Now α_2 is well defined.

Let $c_2(\alpha_2) = \frac{1}{2}c_{12}(\alpha_2) + \frac{1}{2}c_{22}(\alpha_2)$ be the value of the 2-period dual objective with argument α_2, where c_{t2} is the tth term in the sum.

[7] By definition of σ_{i1}, it is easy to see that $\Delta\mathrm{Lr}(\overline{s}|a_1, \sigma_{i1}) = \alpha_{i0}$ does not depend on a_1 because $\Delta\mathrm{Lr}(s|a, \sigma_i)$ does not depend on a, since σ is a feasible asymptotic dual solution.

Simple calculations yield

$$c_{12}(\alpha_2) = \tfrac{1}{2}\sum_{i=1}^{n} u_i(\mu,\alpha_{i1}) = \tfrac{1}{2}\sum_{i=1}^{n} u_i(\mu,\sigma_{i1})\alpha_{i0} = \tfrac{1}{2}\sum_{i=1}^{n}\frac{u_i(\mu,\sigma_{i1})}{\text{Lr}(\overline{s}|a_1,\sigma_{i1})}$$

$$= \tfrac{1}{2}\sum_{i=1}^{n}\frac{u_i(\mu,\sigma_i)}{\text{Lr}(\overline{s}|a_1,\sigma_i)} = \tfrac{1}{2}c_\infty(\mu) \quad \text{for all } a_1\text{, and}$$

$$c_{22}(\alpha_2) = \tfrac{1}{2}\sum_{(i,a_1,b_{i1},s_1,a_2,b_{i2})}\mu(a_1)\Pr(s_1|a_1,b_{i1})\mu(a_2) \quad u_i(a_2,b_{i2})\sigma_{i2}(b_{i2}|a_{i2})\sigma_{i1}(b_{i1}|a_{i1})K(s_1)$$

$$= \tfrac{1}{2}c_\infty(\mu)\sum_{(a_1,s)}\mu(a_1)[\Pr(s|a_1) - \Pr(s|a_1,\alpha_{i1})] = \tfrac{1}{2}c_\infty(\mu),$$

$$\text{where } K(s_1) = \frac{\sum_{a_1}\mu(a_1)[\Pr(s_1|a_1) - \Pr(s_1|a_1,\alpha_{i1})]}{\sum_{a_1}\mu(a_1)\Pr(s_1|a_1,\sigma_{i1}) \quad \text{Lr}(\overline{s}|a_2,\sigma_{i2})}.$$

Therefore, $c_2(\alpha_2) = c_\infty$, so by revealed preference $c_2 \geq c_\infty$. Extending this argument to the t-period dual yields $c_t \geq c_\infty$ for all $t \in \mathbb{N}$. For any (i,t), let $\alpha_{it}(a_i^t, b_i^t, s^{t-1}) = \alpha_{i0}^t \prod_\tau \sigma_{i\tau}(b_{i\tau}|a_{i\tau}, s^{\tau-1})$, where $\alpha_{i0}^t = \quad \text{Lr}(\overline{s}|a_1,\sigma_{i1})$ and

$$\sigma_{i\tau}(b_{i\tau}|a_{i\tau},s^{\tau-1}) = \gamma_i^{\tau-1}(s^{\tau-1})\sigma_i(b_{i\tau}|a_{i\tau}) + (1 - \gamma_i^{\tau-1}(s^{\tau-1}))[a_{i\tau}](b_{i\tau}),$$

and given $\gamma_i^0,\ldots,\gamma_i^{t-1}$, let γ_i^t be de ned recursively by

$$\gamma_i^t(s^t) = \gamma_i^0\prod_{\tau=1}^{t}\frac{\text{Lr}(\overline{s}|a_\tau,\sigma_i) - \text{Lr}(s_\tau|a_\tau,\sigma_i)}{\text{Lr}(\overline{s}|a_\tau,\sigma_i)\text{Lr}(s_\tau|a_\tau,\sigma_{i\tau})}.$$

Now the previous argument for $t = 2$ extends to all $t \in \mathbb{N}$.

The question now arises whether it is possible for c_t to be uniformly bounded strictly above c_∞. For revocable games, the answer is yes. To see this, consider rst the following problem. Given σ, nd vectors $\nu \geq 0$ and $x \geq \underline{x} > -1$ such that

$$\text{Lr}(s|a,\sigma_i) + \sum_{b_i} x_i(a_i,b_i,s)\text{Lr}(s|a,b_i)\sigma_i(b_i|a_i) = \nu(s) \quad \forall(i,a,s).$$

Intuitively, think of $x_i(a_i,b_i,s)$ as coming from $\sum_{b_{i2}} \quad \text{Lr}(s_2|a_2,b_{i2})\sigma_{i2}(b_{i2}|a_i^2,s_1,b_{i1})$ for some σ_{i2} and some (a_2,s_2). By Assumption 1, Pr has full support, therefore $\text{Lr}(s_2|a_2,b_{i2}) \geq \underline{x} > -1$ for some $\underline{x}$. Taking the dual of the above problem yields

$$\max_{\zeta\geq 0,\xi} \; -\sum_{(i,a,s)}\xi_i(a,s) + (1+\underline{x})\sum_{(i,s,a_i,b_i)}\zeta_i(a_i,b_i,s) \quad \text{s.t.}$$

$$\sum_{(i,a)}\xi_i(a,s) \leq 0 \quad \forall s,$$

$$\sum_{a_{-i}}\xi_i(a,s)\text{Lr}(s|a,b_i)\sigma_i(b_i|a_i) = \zeta_i(a_i,b_i,s) \quad \forall(i,a_i,b_i,s),$$

after substituting Lr$(s|a,\sigma_i)$ = Lr$(s|a,\sigma_i) - 1$. The dual above is clearly unbounded if and only if $\mathscr{C}_i(a_i,s)^\circ \cap (-\mathbf{1})^\circ \neq \{\mathbf{0}\}$. This is precisely revocability. In other words, revocability requires existence of some (i,a_i,s) and feasible ξ such that $\sum_{a_{-i}} \xi_i(a,s) <$ 0 or $\sum_{a_{-i}} \xi_i(a,s)$Lr$(s|a,b_i) > 0$ for some b_i in the support of $\sigma_i(a_i)$, or both. (Of course, if $\underline{x} = -1$ then $\sum_{a_{-i}} \xi_i(a,s) < 0$ is necessary for existence of an unbounded dual solution, but $\underline{x} > -1$ by Assumption 1.) It follows that σ_i is detectable.

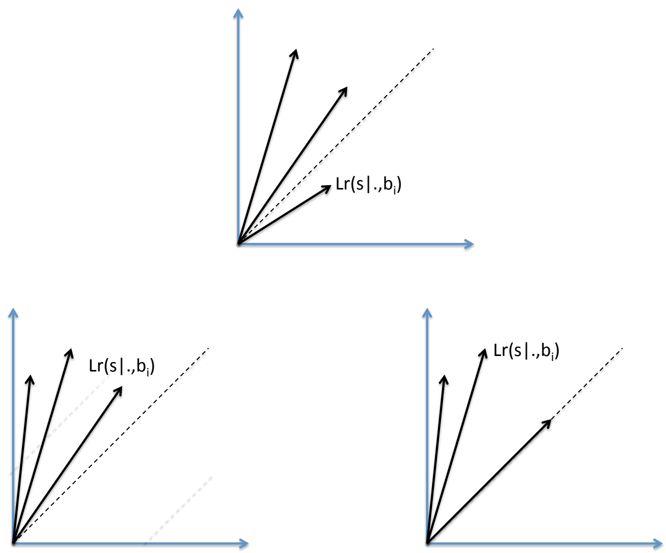

Figure 1: One irregular game (top) and two regular games (bottom). The bottom right game is regular only if $c > -1$.

The top of Figure 1 shows that revocability holds when there is a signal and a player such that the 45 degree line is in the cone generated by likelihood ratios associated with every pure strategy in the support of the player's deviation. On the other hand, the bottom of Figure 1 shows two instances where revocability fails. Intuitively, in these cases, it is impossible for a players to dynamically deviate so that, across periods, they can hamper the mediator's private information regarding players' recommendations.

Now I will show that $c_2 < c_1$ if $c_\infty < c_1$ when the game is irrevocable. De ne the following problem:

$$\min_{\nu \geq 0, \hat{\nu}, x} \nu(s) - \sum_a \mu(a)\hat{\nu}(a,s) \quad \text{s.t.} \quad x \geq \underline{x}, \quad \nu(s) \geq \hat{\nu}(a,s) \quad \forall (a,s),$$

$$\text{Lr}(s|a,\sigma_i) + \sum_{b_i} x_i(a_i,b_i,s)\text{Lr}(s|a,b_i)\sigma_i(b_i|a_i) = \hat{\nu}(a,s) \quad \forall (i,a,s).$$

The dual is given by

$$\chi(s|\sigma) = \max_{\zeta \geq 0, \xi, \kappa} -\sum_{(i,a,s)} \xi_i(a,s) + (1 + \underline{x}) \sum_{(i,s,a_i,b_i)} \zeta_i(a_i,b_i,s) \text{ s.t.}$$

$$\sum_{(i,a)} \xi_i(a,s) = \kappa(a,s) - \mu(a) \quad \forall (a,s), \qquad \sum_a \kappa(a,s) \leq 1 \quad \forall s,$$

$$\sum_{a_{-i}} \xi_i(a,s) \mathrm{Lr}(s|a,b_i) \sigma_i(b_i|a_i) = \zeta_i(a_i,b_i,s) \quad \forall (i,a_i,b_i,s).$$

If the game is irrevocable then $\chi(s|\sigma) > 0$ for some s, which occurs with probability $\pi > 0$.

I will brie y sketch the rest of the proof. First, notice that the set of hampering deviation pro les is a polyhedron. A deviation pro le is extremely unhampering if whenever it is written as a convex combination of deviation pro les, there is zero weight on hampering ones. Hence, there is a bound $\overline{\chi}$ such that $\max_s \chi(s|\sigma) > \overline{\chi} > 0$ for all extremely unhampering σ.

Finally, to show that $c_t \to c_\infty$, let t be su ciently large that $(t-1)\pi\overline{\chi} > 1$. Consider a solution to the dual t-period contractual problem σ^t. If σ_1^t is not hampering then the gain for the t period problem relative to having chosen a hampering pro le is $(c_1 - c_0)/t$. On the other hand, the cost, that is, with probability π the value of the problem is reduced by χ, persists for all subsequent periods. So, there is a cost of $(t-1)\pi\overline{\chi}/t$ from choosing not to hamper the mediator's informational advantage. Since the costs exceed the bene ts, it is optimal to maintain σ_1^t hampering for all blocks longer than t. Now, given this observation, the same logic applies to $\sigma_2^t(s_1)$ by adding one more period to the block, if necessary. Continuing in this way, eventually the average payo s become predominantly arising from hampering strategies, which yields the result for $\lambda \geq 0$.

The case of $\lambda \ngeq 0$ is postponed until the next version.

References

Abreu, D., P. Milgrom, and D. G. Pearce (1990a): \Information and Timing in Repeated Partnerships," *Econometrica*, 59, 1713{33. 5, 11

Abreu, D., D. G. Pearce, and E. Stacchetti (1990b): \Toward a Theory

of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041{1063.

Forges, F. (1986): \An approach to communication equilibria," *Econometrica*, 54, 1375{1385.

Fudenberg, D. and D. Levine (1994): \E ciency and observability with long-run and short-run players," *Journal of Economic Theory*, 62, 103{135. 1, 3, 4, 9, 10

Fudenberg, D., D. Levine, and S. Takahashi (2007): \Perfect public equilibrium when players are patient," *Games and Economic Behavior*, 61, 27{49. 9

Fudenberg, D. and E. Maskin (1986): \The folk theorem in repeated games with discounting or with incomplete information," *Econometrica: Journal of the Econometric Society*, 533{554.

Green, E. and R. Porter (1984): \Noncooperative collusion under imperfect price information," *Econometrica: Journal of the Econometric Society*, 87{100. 3

Kandori, M. and I. Obara (2006): \E ciency in Repeated Games Revisited: The Role of Private Strategies," *Econometrica*, 74, 499{519. 3

Lehrer, E. (1992): \On the Equilibrium Payo s Set of Two Player Repeated Games with Imperfect Monitoring," *International Journal of Game Theory*, 20, 211{226. 13

Myerson, R. (1986): \Multistage games with communication," *Econometrica*, 54, 323{358.

Radner, R., R. Myerson, and E. Maskin (1986): \An Example of a Repeated Partnership Game with Discounting and with Uniformly Ine cient Equilibria," *Review of Economic Studies*, 53, 59{69. 3, 4

Rahman, D. (2008): \But Who Will Monitor the Monitor?" Mimeo. 1, 5, 13

Rahman, D. and I. Obara (2010): \Mediated Partnerships," *Econometrica*, 78, 285{308. 1, 3, 4, 8, 11

Renault, J. and T. Tomala (2004): \Communication equilibrium payo s in repeated games with imperfect monitoring," *Games and Economic Behavior*, 49, 313{344. 11, 13

Wen, Q. (1994): \The" folk theorem" for repeated games with complete information," *Econometrica: Journal of the Econometric Society*, 949{954.