

But Who will Monitor the Monitor?*

David Rahman[†]
University of Minnesota

September 6, 2009.

Abstract

Consider a group of individuals in a strategic environment with moral hazard and adverse selection, and suppose that providing incentives for a given outcome requires a monitor to detect deviations. What about the monitor's deviations? This paper proposes a contractual arrangement that makes the monitor responsible for the monitoring technology (but not the entire firm), and asserts that his deviations are effectively irrelevant. Hence, nobody needs to monitor the monitor. The contract successfully provides incentives even when the monitor's observations are not only private, but costly, too. We also characterize exactly when such a contract can provide monitors with the right incentives to perform. In doing so, virtual enforcement is characterized.

JEL Classification: D21, D23, D82.

Keywords: contracts, private monitoring, communication equilibrium, costly subjective evaluation.

*This paper subsumes [Rahman \(2005a\)](#), which in turn is based on Chapter 3 of [Rahman \(2005b\)](#).

[†]Financial support from the Spanish Ministry of Education's Research Grant No. SEJ 2004-07861 while at Universidad Carlos III de Madrid is gratefully acknowledged.

1 Introduction

Ann owns a restaurant. She hires Bob to tally the till every night and report back any mismatch between the till and that night's bills. Ann is too busy to check the till herself and has to trust what Bob says. How can Ann provide Bob with appropriate incentives to exert the effort required to tally the till and report back the truth?

Ann's problem, basic as it is, seems to have eluded systematic analysis by economists. In studying incentives, most economists have focused on output-contingent contracts, such as bonuses for sales reps.¹ Thus, a great way of convincing a salesperson to exert effort is to promise him or her a greater reward the more he or she sells. However, this kind of contract gives Bob perverse incentives, since only he can know if there is a mismatch between the till and the bills. Hence, if Ann paid Bob a bonus for reporting a mismatch then Bob would just report it without tallying the till, and similarly if the bonus was for reporting no mismatch. Some economists have suggested ways to provide incentives for truth-telling,² which in this setting boils down to simply paying Bob the same amount regardless of what he says to make him indifferent between honesty and deception. However, this contract cannot help Ann either because then nothing would prevent Bob from neglecting to tally the till.

This kind of problem is pervasive. For instance, consider TSA airport inspectors that sit behind an X-ray machine, watching suitcases pass them by. Their "output" is paying attention—only they know if they are studying the suitcases in front of them or just daydreaming. A related example is under-aged drinking. Without the right incentives, most bartenders would naturally prefer not having to check IDs before serving alcoholic drinks to their customers. Finally, in a classroom, teachers usually seek to provide students with the right incentives to study, but this is unobservable.

¹A classic example is Holmström (1982), but see also Legros and Matsushima (1991), Legros and Matthews (1993), Strausz (1997) and d'Aspremont and Gérard-Varet (1998).

²See the literature on subjective evaluation, especially the work by Prendergast (1999), Levin (2003), MacLeod (2003) and Fuchs (2007). In a principal-agent model where only the principal observes output (i.e., subjective evaluation), they argue that the principal must be indifferent over reports to tell the truth. However, they all assume that subjective evaluations are costless, and their contract breaks down if observing output is costly—no matter how small this cost. In this paper, we accommodate *costly* subjective evaluations by providing incentives for reporting accuracy.

I propose the following solution to Ann’s problem: Ann can motivate Bob to exert effort and report truthfully by sometimes secretly taking money from the till herself and offering him the following deal: if Ann took some money, she will pay Bob his wage only when he reports a mismatch; if Ann did not take any money, she will pay Bob only when a mismatch is not reported. Bob’s incentives are now aligned with Ann’s. Indeed, if Bob doesn’t bother tallying the till, he won’t know what to tell Ann in order to make sure he gets paid. On the other hand, if he does his job he’ll discover whether or not there is a mismatch and deduce whether or not Ann took some money. Only after tallying the till will Bob know what to tell Ann in order to receive his wage. By offering Bob this perhaps “contrived” contract, Ann can now rest assured that Bob will be honest and obedient.

Contrived though it may seem, this kind of contract is ubiquitous. Indeed, the TSA uses a version they call “covert testing” to evaluate inspectors (TSA, 2004, page 5). Police officers go undercover to bars asking for alcoholic drinks to make sure that bartenders check IDs (Cheslow, 2005). Finally, of course, students are given the right incentives by being tested. As every teacher knows, a good test question is one which a student can only answer if he or she understands the course material.

The insight behind Bob’s contract has far-reaching consequences for understanding the role of monitoring in organizations—exploring them is the purpose of this paper. Since Alchian and Demsetz (1972) posed the basic question of how to remunerate monitors,³ it has generated much academic debate. This paper adds to the debate by constructing a theoretical model whose answer is that nobody needs to monitor the monitor if he is made responsible for the monitoring technology. The model also describes just how to make the monitor responsible with a more general version of Bob’s contract, and characterizes exactly when this contract is enforceable.

Unlike the rest of the literature, this paper accommodates *costly private monitoring*. Previously, a monitor’s observations were assumed to be either publicly verifiable (Footnote 1) or costless (Footnote 2), and all the solutions proposed thus far fail to provide the right incentives in this richer environment. For instance, the literature on subjective evaluation aptly argues that a costless private monitor must be indifferent over his reports, otherwise he will have the incentive to lie. However, making the monitor indifferent discourages observing output at a cost—no matter how small.

³Juvenal asked a very similar question when he argued that no husband can trust his wife to be faithful by having someone guard her to guarantee her celibacy while he is away (en.wikipedia.org/wiki/quis_custodiet_ipsos_custodes). But see also en.wikipedia.org/wiki/eunuch.

I begin my analysis in [Section 2](#) by applying Bob’s contract to a firm with two agents: a worker (Friday) and a costly private monitor (Robinson). I show how to provide the monitor with the right incentives to perform at the cost of occasional shirking by the worker. The principal overcomes the monitor’s apparent informational advantage by allocating different private information (incentive compatible effort recommendations) to different agents, and making the monitor’s reward contingent on information allocated to the worker. Formally, I consider contracts that form a *communication equilibrium* ([Forges, 1986](#); [Myerson, 1986](#)).

Most of the time the owner asks Friday to work but once in a while he secretly asks Friday to shirk. In the former case, he pays Friday only if Robinson verifies his effort. In the latter, Friday gets nothing. Robinson is rewarded as follows: if Friday was asked to work then Robinson will be paid only if he reports back that Friday worked, whereas if Friday was asked to shirk then Robinson will be paid only if he reports back that Friday shirked. This contract rewards Robinson for reporting accuracy, since now Robinson has the incentive to acquire costly information (i.e., monitor) and reveal it. Indeed, if Robinson shirks he won’t know what to report in order to get paid, whereas if he monitors then he’ll observe Friday’s behavior and thereby deduce the owner’s effort recommendation, which is just what he wants to secure payment. Therefore, every agent is honest and obedient in equilibrium.

For the next main result of this paper, I extrapolate from the previous example and consider a general contracting environment. I reconcile the following infinite regress inherent in monitoring. Suppose that providing workers with incentives to exert effort requires costly private monitoring to detect their deviations. What about the monitor’s deviations? [Theorem 3](#) asserts that the monitor’s deviations are effectively irrelevant. Indeed, if they are detectable then they can be discouraged with contingent payments similar to Bob’s. Otherwise, if they are undetectable then the deviations themselves still detect workers’ deviations by virtue of being undetectable, so they continue to fulfill the intended monitoring role. Evidently, this argument also applies to the monitor’s deviations from these deviations, and so forth. Technically, this infinite regress is reconciled by showing that under reasonable assumptions (e.g., if every agent has finitely many choices) not every behavior by the monitor can have a profitable, undetectable deviation. Therefore, to induce workers’ effort with infrequent monitoring, workers’ deviations must be detectable with some monitoring behavior, but deviations away from the monitoring behavior itself need not be detectable. Heuristically, *nobody needs to monitor the monitor*.

2 Example

Example 1 (*Robinson and Friday*). Consider a principal and two risk neutral agents, Robinson (the row player) and Friday (the column player), who interact with payoffs in the left bi-matrix below. Intuitively, Friday is a worker and Robinson is a monitor. Each agent’s effort is costly—with cost normalized to unity—but unobservable.

	work	shirk
monitor	0, 0	0, 1
rest	1, 0	1, 1

Utility Payoffs

	work	shirk
monitor	1, 0	0, 1
rest	1/2, 1/2	1/2, 1/2

Signal Probabilities

After actions have been taken, Robinson privately observes one of two possible signals, g and b . Their conditional probability (or *monitoring technology*) appears in the right bi-matrix above. In words, if Robinson monitors he observes Friday’s effort, but if he rests then his observation is completely uninformative.⁴ Finally, after Robinson observes the realized signal, he makes a verifiable report to the principal.

The principal wants Friday to work. If monitoring were costless then—following the literature on subjective evaluation ([Footnote 2](#))—the principal could enforce the action profile (monitor,work) by paying Robinson a wage independent of his report. Under this contract, Robinson would be happy to monitor and report truthfully, and Friday could therefore be rewarded contingent on Robinson’s verifiable report.

With costly monitoring, Robinson’s effort becomes an issue. Suppose that the principal wants to enforce (rest,work) on the grounds that monitoring is unproductive. On the one hand, this is impossible, since if Robinson rests Friday’s expected payment cannot depend on his own effort, so he will shirk. On the other, if Robinson’s signals are publicly verifiable then not only can the principal easily enforce (monitor,work), but he can also *virtually enforce* (rest,work)—i.e., enforce an outcome arbitrarily close to it—using [Holmström’s](#) group penalties. Intuitively, if news is good everyone gets paid and if news is bad nobody gets paid. Specifically, the principal can induce Friday to always work and Robinson to secretly monitor with small but positive probability σ by paying Robinson \$2 and Friday $\$1/\sigma$ if g and both agents zero if b .

⁴Alternatively, we may have assumed that if Robinson rests he observes “no news.” Our current assumption helps to compare with the literature that relies on publicly verifiable monitoring.

If Robinson's costly observations are unverifiable, [Holmström's](#) contracts break down, since Robinson will then just report g and rest, so Friday will shirk. Furthermore, even though Robinson would happily tell the truth with a wage independent of his report, he would never monitor, and again Friday would shirk. This begs the question: How can we get Friday to work when Robinson's signal is costly and private?

Having Friday always work is impossible, since then Robinson will never monitor, so Friday will shirk. However, the principal can virtually enforce (rest,work) by asking Friday to shirk occasionally and secretly correlating Robinson's payment with Friday's recommendation, thereby "monitoring the monitor." Indeed, the following contract is incentive compatible given $\mu \in (0, 1)$ and $\sigma \in (0, 1]$: (i) Robinson is told to monitor with probability σ (and rest with probability $1 - \sigma$), (ii) Friday is independently told to work with probability μ (to shirk with $1 - \mu$), and (iii) the principal enforces the following payments to Robinson and Friday, respectively, contingent on his recommendations and Robinson's report:

	(monitor,work)	(monitor,shirk)	(rest,work)	(rest,shirk)
g	$1/\mu, 1/\sigma$	$0, 0$	$0, 0$	$0, 0$
b	$0, 0$	$1/(1 - \mu), 0$	$0, 0$	$0, 0$

Thus, Friday is paid with [Holmström's](#) contract, whereas Robinson is paid $\$1/\mu$ if he reports g when (monitor,work) was recommended and $\$1/(1 - \mu)$ if he reports b when (monitor,shirk) was recommended. He is not told Friday's recommendation—this he must discover by monitoring. Clearly, Friday has the incentive to obey the principal's recommendations if Robinson is honest and obedient. To see that Robinson will abide by the principal's requests under this contract, suppose that he was asked to monitor. If he does monitor, then clearly it is optimal for him to be honest, with expected payoff of $\mu(1/\mu) + (1 - \mu)[1/(1 - \mu)] = 2$. After resting instead, his expected payoff equals $1 + \mu(1/\mu) = 2$ if he reports g , and $1 + (1 - \mu)[1/(1 - \mu)] = 2$ if he reports b .

As $\sigma \rightarrow 0$ and $\mu \rightarrow 1$, Robinson and Friday's behavior under the above contract tends to the profile (rest,work) with payments that make the behavior incentive compatible along the way. In other words, (rest,work) is *virtually enforceable*. This requires arbitrarily large payments. In reality, feasible payments may be bounded. Nevertheless, virtual enforcement is still a useful benchmark for attainable outcomes. On the one hand, it describes what is attainable with sufficiently large payments. On the other, interpreting payments as continuation values in a dynamic game, virtual enforcement describes what is attainable as agents become unboundedly patient.

Although (rest,work) is virtually enforceable, Friday shirks with positive probability along the way. In the case of public monitoring, (rest,work) is virtually enforceable by incurring the cost of monitoring Friday (Robinson’s effort) with small probability. With private monitoring, an additional cost must be incurred (also with small probability): the cost of monitoring Robinson. This cost is captured by the foregone productivity resulting from Friday shirking.

Robinson’s contract pays him for matching his report to Friday’s recommendation—he faces a “trick question” whose answer the principal already knows. This way, Robinson is rewarded for reporting accuracy: he is responsible for the monitoring technology, and not for any output that Friday’s effort might generate. As such, he must not observe Friday’s recommendation, since his job is only to confirm it to the principal. Therefore, a problem with this contract is that it is not robust to “collusion:” both agents could avoid effort if Friday told Robinson his recommendation. However, this is cheap talk—it still is an equilibrium that they don’t share this information. (Section 4.3 discusses collusion in more detail.)

Alchian and Demsetz argued for making Robinson the principal. However, if Robinson was the principal then he would never verify Friday’s effort, as Friday’s payment would come from his own pocket *after* Friday’s effort had already been exerted. This argument relies on the fact that Robinson and Friday will not meet in the future, so that Friday cannot threaten to retaliate Robinson if he “cheats.”⁵ In addition, Robinson must not be telling people what to do (giving secret recommendations of effort), because otherwise the above contracts would break down.

If Friday was the principal it would also be impossible to provide the right incentives for two reasons. Firstly, if Robinson was the only one who could verify Friday’s output at a cost then Friday would have to ask the trick questions to Robinson himself. In this case, it would be optimal for him to disobey his own recommendation to himself in order to save paying Robinson his wage. Secondly, it would be impossible to save on the costs of monitoring Friday by having Robinson monitor randomly if Friday was the one telling Robinson when to monitor.

If recommendations are not verifiable, it is still possible to virtually enforce (rest,work) without a third party by asking Friday if he worked. See Section 5.2 for the details.

⁵Several authors have “used time” to solve the principal-agent problem, such as Levin (2003) and Fuchs (2007). However, for any fixed discount factor less than one, there is always a residual incentive problem that cannot be solved dynamically.

3 Model

Let $I = \{1, \dots, n\}$ be a finite set of risk neutral agents, A_i a finite set of actions available to any agent $i \in I$, and $A = \prod_i A_i$ the (nonempty) space of action profiles. Let $v_i(a)$ denote the utility to agent i from action profile $a \in A$. A *correlated strategy* is a probability measure $\mu \in \Delta(A)$.⁶ Let S_i be a finite set of private signals observable only by agent $i \in I$ and S_0 a finite set of publicly verifiable signals. Let $S = \prod_{j=0}^n S_j$ be the (nonempty) product space of all observable signals. A *monitoring technology* is a measure-valued map $\Pr : A \rightarrow \Delta(S)$, where $\Pr(s|a)$ stands for the conditional probability that $s \in S$ was observed given that $a \in A$ was played.

Incentives are provided with linear transfers. An *incentive scheme* is any function $\zeta : I \times A \times S \rightarrow \mathbb{R}$ that assigns individual payments contingent on *recommended* actions and *reported* signals. Recommendations and reports are assumed verifiable.

Time elapses as follows. Firstly, agents agree on a *contract* (μ, ζ) . A profile of suggestions is drawn according to μ and made to agents confidentially and verifiably by the principal.⁷ Agents now simultaneously take unverifiable and unobservable actions. Next, agents observe their private signals and submit a report before a public signal realizes (the order of signals is not essential, just simplifying). Finally, the principal pays agents according to ζ contingent on recommendations and reports.

If every agent obeys his recommendation and reports truthfully, the expected utility to agent i from a given contract (μ, ζ) equals

$$U_i(\mu, \zeta) = \sum_{a \in A} \mu(a) v_i(a) - \sum_{(a, s)} \mu(a) \zeta_i(a, s) \Pr(s|a).$$

Of course, agent i may disobey his recommendation and lie about his private signal. A *reporting strategy* is a map $\rho_i : S_i \rightarrow S_i$, where $\rho_i(s_i)$ is the reported signal when agent i privately observes s_i . Let R_i be the set of i 's reporting strategies. The *truthful reporting strategy* is the identity map $\tau_i : S_i \rightarrow S_i$ with $\tau_i(s_i) = s_i$. For every agent i and pair $(b_i, \rho_i) \in A_i \times R_i$, the probability that $s \in S$ is reported if everyone else is honest and plays $a_{-i} \in A_{-i}$ equals

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

⁶If X is a finite set, $\Delta(X) = \{\mu \in \mathbb{R}_+^X : \sum_x \mu(x) = 1\}$ is the set of probability vectors on X .

⁷See Section 5.2 for a contract modification if recommendations are not verifiable.

A contract (μ, ζ) is called *incentive compatible* if honesty and obedience is optimal:

$$\sum_{a_{-i}} \mu(a) [v_i(a_{-i}, b_i) - v_i(a)] \leq \sum_{(a_{-i}, s)} \mu(a) \zeta_i(a, s) [\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a)] \quad (*)$$

for every (i, a_i, b_i, ρ_i) . In other words, (μ, ζ) is incentive compatible if μ is a *communication equilibrium* (Myerson, 1986; Forges, 1986) of the game induced by ζ .⁸

Definition 1. A correlated strategy μ is *exactly enforceable* (or simply *enforceable*) if an incentive scheme ζ exists such that (μ, ζ) is incentive compatible. Call μ *virtually enforceable* if a sequence $\{\mu^m\}$ of enforceable correlated strategies exists with $\mu^m \rightarrow \mu$.

A *strategy* for agent i is a map $\sigma_i : A_i \rightarrow \Delta(A_i \times R_i)$, where $\sigma_i(b_i, \rho_i|a_i)$ is the probability that i plays (b_i, ρ_i) when recommended a_i . Let $\Pr(\mu)$ be the vector of report probabilities if everyone is honest and obedient, defined by $\Pr(s|\mu) = \sum_a \mu(a) \Pr(s|a)$ for each s . Let $\Pr(\mu, \sigma_i)$ be the vector of report probabilities if agent i plays σ_i instead of honesty and obedience, defined by

$$\Pr(s|\mu, \sigma_i) = \sum_{a \in A} \mu(a) \sum_{(b_i, \rho_i)} \Pr(s|a_{-i}, b_i, \rho_i) \sigma_i(b_i, \rho_i|a_i)$$

for each signal profile s .

Definition 2. Given any subset of action profiles $B \subset A$, a strategy σ_i is called *B-detectable* if $\Pr(s|a) \neq \Pr(s|a, \sigma_i)$ for some $a \in B$ and $s \in S$.⁹ Otherwise, σ_i is called *B-undetectable*. A strategy is simply *detectable* if it is *A-detectable*, etc.

Intuitively, a strategy is *B-detectable* if there is a profile of recommendations in B such that the report probabilities it induces differs from that induced by honesty and obedience, assuming others are honest and obedient.

We begin with an intuitive characterization of enforceable outcomes. For any correlated strategy μ , consider the following zero-sum two-person game between the principal and a “surrogate” for the agents. The principal chooses an incentive scheme ζ and the surrogate chooses a strategy σ_i for some agent i . The principal pays the surrogate the expected deviation gains from i playing σ_i instead of being honest and obedient, $\sum_{(a, b_i, \rho_i)} \mu(a) \sigma_i(b_i, \rho_i|a_i) [(v_i(a_{-i}, b_i) - v_i(a)) - \sum_s \zeta_i(a, s) (\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a))]$.

⁸Strictly speaking, the communication equilibrium is defined with the principal being a disinterested player who chooses payments to agents contingent on a mediator’s recommendations.

⁹We abuse notation by identifying Dirac measure $[a] \in \Delta(A)$ with the action profile $a \in A$.

By construction, μ is enforceable if and only if the value of this game is zero. Notice that this value is at least zero for the surrogate, since he could always have his agents play honest and obediently. Therefore, the value is zero if and only if there is an incentive scheme that discourages every strategy by the surrogate, i.e., that makes every strategy unprofitable relative to honesty and obedience. By the Minimax Theorem, the value of the game is independent of the order of moves. Hence, μ is enforceable if and only if for every strategy there is an incentive scheme that discourages it, where different schemes may be used to discourage different strategies. Intuitively, for μ to be enforceable it suffices that the principal can discourage strategies one by one.

Pick any strategy σ_i . If it is $\text{supp } \mu$ -detectable¹⁰ then there exists $a \in \text{supp } \mu$ such that $\Pr(a) \neq \Pr(a, \sigma_i)$. Hence, there are signals whose probability increases with σ_i (“bad” news) and others whose probability decreases (“good” news). The following incentive scheme discourages σ_i : choose a sufficiently large wedge between good and bad news after a is recommended such that the monetary loss outweighs any utility gain from playing σ_i . On the other hand, if σ_i is $\text{supp } \mu$ -undetectable then the surrogate’s payoff is unaffected by the incentive scheme. Hence, if σ_i gives the surrogate a positive payoff then there is nothing the principal can do to discourage it.

Theorem 1 (Minimax Lemma). *A correlated strategy μ is enforceable if and only if every $\text{supp } \mu$ -undetectable strategy σ_i is μ -unprofitable, i.e.,*

$$\Delta v_i(\mu, \sigma_i) = \sum_{(a, b_i, \rho_i)} \mu(a) \sigma_i(b_i, \rho_i | a_i) [v_i(a_{-i}, b_i) - v_i(a)] \leq 0.$$

In principle, to verify that μ is enforceable one must find an incentive scheme and check that every strategy is unprofitable. By the Minimax Lemma it is enough to assume that $\zeta \equiv 0$ and only verify that every undetectable deviation is unprofitable. As a result, if every relevant strategy is $\text{supp } \mu$ -detectable, then the consequent of the Minimax Lemma holds vacuously, and μ is enforceable regardless of the utility profile $v : I \times A \rightarrow \mathbb{R}$. What makes a strategy relevant? Clearly, only strategies that differ from honesty and obedience with positive probability are relevant. Furthermore, since reports are costless in terms of utility, only actions that differ from recommendations matter. This intuition leads to the following definition and result.

Definition 3. Given $B \subset A$, a strategy σ_i is called a *B-disobedience* if $\sigma_i(b_i, \rho_i | a_i) > 0$ for some $a_i \in B_i$ and $b_i \neq a_i$, where $B_i = \{b_i \in A_i : \exists b_{-i} \in A_{-i} \text{ s.t. } b \in B\}$ is the projection of B on A_i . An *A-disobedience* is called simply a *disobedience*.

¹⁰Let $\text{supp } \mu = \{a \in A : \mu(a) > 0\}$ be the set of action profiles with positive probability under μ .

Theorem 2. *Fix any arbitrary correlated strategy μ . Every supp μ -disobedience is supp μ -detectable if and only if for any profile of utility functions, μ is enforceable.*

For every disobedience to be detectable, different action profiles may be used to detect different disobediences. This key feature renders such a requirement much weaker than other conditions in the literature, such as individual full rank (IFR) by Fudenberg et al. (1994). To illustrate, consider a monitoring technology such that every disobedience is detectable but IFR fails. In fact it even fails to satisfy *local* IFR of d’Aspremont and Gérard-Varet (1998).

Example 2. There are two publicly verifiable signals, $S = S_0 = \{x, y\}$, and two agents, Ann and Bob. Ann has two choices, $\{U, D\}$, and Bob has three, $\{L, M, R\}$. The monitoring technology Pr is given in the bi-matrix below.

	L	M	R
U	1, 0	0, 1	1/2, 1/2
D	1, 0	0, 1	1/3, 2/3

If Ann plays U and Bob plays $\frac{1}{2}[L] + \frac{1}{2}[M]$ then the probability over signals equals what it would have been had Bob played R . Similarly, if Ann plays D then Bob can deviate from R to play $\frac{2}{3}[L] + \frac{1}{3}[M]$ without changing the probability over signals. It is therefore impossible to even virtually enforce R with transfers contingent only on signals if Bob strictly prefers playing L and M , since there always exists a profitable deviation without monetary loss. However, every disobedience is detectable, because for any deviation by Bob there is a mixed strategy by Ann that detects it. By correlating Bob’s payment with Ann’s recommendation, the principal can keep Bob from knowing the proportion with which he ought to mix between L and M for his payment to equal what he would obtain by playing R . This renders R enforceable.

Notice that only the support of a correlated strategy appears in Theorem 2 above. Intuitively, this is because payments are recommendation-contingent, so only the set of action profiles with positive probability matters for enforcement.

Corollary 1. *Every B -disobedience is B -detectable if and only if for any profile of utility functions, every correlated strategy with support equal to B is enforceable.*

By Corollary 1, every disobedience is detectable if and only if every completely mixed correlated strategy is enforceable. Approaching an arbitrary correlated strategy with completely mixed ones, it becomes virtually enforceable. The converse is also true.

Proposition 1. *Every disobedience is detectable if and only if for any profile of utility functions, every correlated strategy is virtually enforceable.*

Our next result characterizes virtual enforcement of a fixed correlated strategy, rather than every one. To this end, fix a correlated strategy μ with support $B \subset A$. By [Corollary 1](#), if every C -disobedience is C -detectable for some $C \supset B$ then for any utility profile, μ is virtually enforceable. Indeed, since C contains B , μ is approachable with correlated strategies whose support equals C , and every such correlated strategy is enforceable. However, μ can be virtually enforceable for every utility profile even if this requirement fails. To see this, consider an example.

Example 3. Two agents, two public signals, the following monitoring technology:

	L	M	R
U	1, 0	1, 0	1, 0
D	1, 0	0, 1	0, 1

Clearly, (U, L) is not enforceable for every profile of utility functions, since there exist $\{(U, L)\}$ -undetectable $\{(U, L)\}$ -disobediences such as playing D if asked to play U . It is also easy to see there exists a C -undetectable C -disobedience for every $C \supset \{(U, L)\}$. However, (U, L) is virtually enforceable, since either $[(D, M)]$ or $[(D, R)]$ can be used to detect $\{(U, L)\}$ -disobedient deviations. The key condition here is that every $\{(U, L)\}$ -disobedient deviation plan is detectable.

Theorem 3. *Fix any correlated strategy μ . Every supp μ -disobedience is detectable if and only if for any profile of utility functions, μ is virtually enforceable.*

[Theorem 3](#) is one of the main results of the paper. It shows that μ is virtually enforceable for every utility profile as long as every disobedience from μ is detectable with some infrequent behavior—call it “monitoring.” Crucially, there is no requirement on disobediences to behavior outside of μ , *even monitoring*. In other words, deviations from monitoring need not be detectable.

To make intuitive sense of all this, let $B \subset A$ be the support of μ . Recall that by the Minimax Lemma, we may discourage disobediences one by one. Suppose that, to detect a disobedience $\sigma_i(a_i)$ away from $a_i \in B_i$, some $a_j \notin B_j$ must be played infrequently by $j \neq i$. Call this “monitoring.” What if a_j itself has a profitable deviation $\sigma_j(a_j) \in \Delta(A_j \times R_j)$? After all, the condition of [Theorem 3](#) purposely says nothing about detection outside B .

If such $\sigma_j(a_j)$ is detectable then it can be discouraged by some incentive scheme. If on the other hand $\sigma_j(a_j)$ is undetectable then playing $\sigma_j(a_j)$ instead of a_j still detects deviations from a_i by virtue of being undetectable, in other words, *it's still monitoring*. Similarly, undetectable deviations from $\sigma_j(a_j)$ detect deviations from a_i , and so on. Proceeding iteratively, since the game is finite there must be detecting behavior without a profitable, undetectable deviation.

This intuitive argument completes our answer to the question “But who will monitor the monitor?” The principal monitors the monitor’s detectable deviations at the cost of his workers shirking infrequently, and nobody needs to monitor the monitor’s undetectable deviations. This is accomplished with a contract that aligns the monitor’s incentives with the principal’s by making the monitor responsible for the monitoring technology. Finally, the monitor is made responsible with contractual terms that follow Robinson’s incentive scheme in [Example 1](#). These terms provide monitors with incentives for reporting accuracy, especially when monitoring is costly and private.

Since [Theorem 3](#) only depends on the support of μ , we obtain an immediate corollary.

Corollary 2. *Fix a subset $B \subset A$. Every B -disobedience is detectable if and only if for any profile of utility functions, every correlated strategy with support contained in B is virtually enforceable.*

4 Discussion

The results above are useful for understanding the requirements on a monitoring technology that allow for the right incentives to be provided to agents. This section comments on possible extensions and limitations of the model and results.

We begin this section by extending [Theorem 3](#) to characterize virtual enforcement for a fixed profile of utility functions. This extension is more subtle than [Theorems 1](#) and [3](#) would suggest, and delivers new economic insights. We continue by characterizing the contractual added value of considering recommendation-contingent payments and noting that [Theorem 3](#) relies on them. Next, we discuss collusion and characterize contracts that dissuade multilateral deviations. Finally, we study genericity. Specifically, we derive weak sufficient conditions under which the set of monitoring technologies such that every disobedience is detectable is generic.

4.1 Fixed Utility Functions

Below we characterize virtual enforcement for a fixed utility profile. Although exact enforcement has a simple characterization ([Proposition 1](#)), a corresponding result for virtual enforcement is trickier to obtain. To see this, notice that on the one hand virtually enforcing an outcome μ does not require that every supp μ -disobedience be detectable. For instance, an unprofitable disobedience may be detectable and yet μ may still be virtually enforceable. On the other hand, it is not enough that every profitable supp μ -disobedience be detectable, as [Example 4](#) below shows.

Example 4. Consider the following variation on Robinson and Friday ([Section 2](#)).

	work	shirk	solitaire		work	shirk	solitaire
monitor	0, 0	0, 1	0, 1	monitor	1, 0	0, 1	1, 0
rest	0, 0	0, 1	0, 0	rest	1/2, 1/2	1/2, 1/2	1/2, 1/2

Utility Payoffs
Signal Probabilities

Assume that the signal is *publicly verifiable* and Robinson's utility is constant. Clearly, the profile (rest,work) is not enforceable because a deviation by Friday to shirk is (rest,work)-profitable and $\{(rest,work)\}$ -undetectable. Moreover, (rest,work) is *not virtually enforceable* either. Indeed, for Friday to ever work it is clear that Robinson must monitor with positive probability. But then no contract can discourage Friday from playing solitaire instead of working, since playing solitaire when asked to work is undetectable and weakly dominant. On the other hand, every (rest,work)-profitable disobedience is detectable, because a (rest,work)-profitable strategy must involve shirking with positive probability and shirking is detectable.

The problem here is that solitaire weakly dominates working and they are indistinguishable. Clearly, if solitaire strictly dominated working there would exist a (rest,work)-profitable, undetectable strategy, rendering (rest,work) virtually unenforceable. Moreover, if Friday's payoff from (rest,solitaire) was negative instead of zero then (rest,work) would be virtually enforceable because playing solitaire when asked to work would be unprofitable if Robinson monitored with low probability.

Removing solitaire restores virtual enforcement of (rest,work). This takes place not because every (rest,work)-profitable deviation is detectable (it is true with or without solitaire), but because it is *uniformly detectable*, i.e., the utility gains from every (rest,work)-profitable disobedience can be uniformly outweighed by monetary losses.

To understand what it takes to virtually enforce an outcome, we will characterize it intuitively as follows: a correlated strategy is virtually enforceable if and only if every profitable disobedience is uniformly and credibly detectable.

To describe “uniform detection,” given an enforceable correlated strategy μ , we now ask how large transfers must be to enforce it. To this end, let us introduce some notation. For any strategy σ_i and any correlated strategy μ , write

$$\|\Delta \Pr(\mu, \sigma_i)\| = \sum_{s \in S} \left| \sum_{(a, b_i, \rho_i)} \mu(a) [\sigma_i(b_i, \rho_i | a_i) \Pr(s | a_{-i}, b_i, \rho_i) - \Pr(s | a)] \right|.$$

Intuitively, this norm describes the statistical difference between abiding by μ and deviating to σ_i . Thus, σ_i is μ -undetectable if and only if $\|\Delta \Pr(\mu, \sigma_i)\| = 0$.

Theorem 4. *A correlated strategy μ is virtually enforceable if and only if there exists $z \geq 0$ such that every μ -profitable strategy σ_i is detectable by some correlated strategy η for which both*

- (i) $\Delta v_i(\eta, \sigma_i) < z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_i)\|$ and
- (ii) $\Delta v_j(\eta, \sigma_j) \leq z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_j)\|$ for every other agent j and strategy σ_j .

Intuitively, Theorem 4 says that to virtually enforce a correlated strategy, it is both necessary and sufficient that all its profitable deviations be discouraged (i) uniformly and (ii) credibly. As usual, different actions may be used to detect different deviations. Formally, uniform detection means that for the same fixed z , every strategy σ_i must impact the magnitude of z -weighted probabilistic changes enough to outweigh its deviation gains. Therefore, transfers bounded within z can provide incentives against all μ -profitable deviations, perhaps with different $\eta \in \Delta(A)$ for different σ_i .

To explain the need for credibility, compare Theorem 4 above with Theorem 3, where “credible monitoring” is unnecessary. There, *every* disobedience is potentially profitable, so ought to be detectable. Here, with fixed utility functions, even if some disobedience σ_i is undetectable, it may nonetheless be discouraged with behavior η by others that makes σ_i unprofitable (as in a correlated equilibrium without transfers), rather than by using contingent money payments. However, if this specific behavior is not credible then there may exist a η -profitable strategy σ_j by some other agent such that σ_i becomes profitable once again given η and σ_j .¹¹

¹¹To see that credibility matters, simply add a row to the table in Example 4 above with utility payoffs

-1, 0	-1, 1	-1, 0
-------	-------	-------

 and signal probabilities

1, 0	0, 1	1, 0
------	------	------

. Now there is an action for Robinson that is strictly dominated and indistinguishable from monitoring, yet uniformly detects all of Friday’s (rest, work)-profitable deviations.

4.2 The Value of Mediated Contracts

The results of [Sections 2 and 3](#) crucially rely on incentive schemes that depend on the principal’s recommendations. Obviously, such schemes yield a weak improvement for the principal—as a result of acquiring a richer contract space—relative to schemes that just depend on reported signals. [Examples 1 and 2](#) show that such schemes can yield a strict improvement, too. In this subsection we suggest a way of interpreting this contractual enrichment. Intuitively, we argue that recommendation-contingent schemes allow monitors to effectively behave as auditors, i.e., to monitor “after the fact.” This enhances the role of monitoring. To this end, fix a correlated strategy μ . A strategy σ_i is called *detectable at μ* if $\Pr(s|\mu) \neq \Pr(s|\mu, \sigma_i)$ for some $s \in S$.

Proposition 2. *Fix any correlated strategy μ . Every strategy that is undetectable at μ is also μ -unprofitable if and only if μ is enforceable with an incentive scheme that does not depend on the principal’s recommendations.*

It easily follows from this result that every μ -disobedient strategy is detectable at μ if and only if for any profile of utility functions, μ is enforceable with an incentive scheme that is independent of recommendations. [Proposition 2](#) captures the value-added of mediated contracts relative to incentive schemes that do not depend on the principal’s recommendations. This value is captured by the difference between μ -detectability and detectability at μ .

To interpret this difference, fix a correlated strategy μ consider a hypothetical game of hide and seek between the principal and a surrogate for the agents. The surrogate chooses a disobedience and the principal chooses an action profile in the support of μ . If for any disobedience the principal can react and find an action profile that detects it then the principal wins and μ is enforceable for any profile of utility functions. In other words, this is as if the principal chooses a correlated strategy *after* agents choose strategies, in order to detect them. To illustrate, recall [Example 1](#). Suppose that Robinson is asked to monitor but instead chooses to rest and report g . The principal can “react” by asking Friday to shirk, which would lead to b if Robinson monitored and reported truthfully. Similarly, if Robinson plans to rest and report b then Friday can be asked to work instead, and Robinson’s deviation is detected again.

On the other hand, if the principal moves first and subsequently the surrogate chooses a disobedience then he may be able to find one that is undetectable at μ . Thus, Robinson monitoring is not enforceable without recommendation-contingent payments.

As an illustration, consider enforcing a pure-strategy profile a . By [Theorem 1](#), this requires that every a -profitable disobedience be $\text{supp } [a]$ -detectable. In this case, $\text{supp } [a]$ -detectability coincides with detectability at a . Since agents receive only one recommendation under a , there is no use for mediated contracts, so by [Proposition 2](#) detectability at a characterizes enforcement with mediated contracts as well as with schemes independent of recommendations. However, when enforcing a correlated strategy with non-singleton support, the two contract spaces differ, and as a result so do the appropriate notions of detectability.

This leads to the observation that [Theorem 3](#) relies on mediated contracts. Indeed, when considering a monitor's undetectable deviations, we did not refer to the particular correlated strategy that was being enforced, since we were focusing on virtual enforcement. As a result, without such contracts the result would not follow.

4.3 Coalitional Deviations

A notable weakness of secret contracts is not being collusion-proof. To illustrate, in our leading example ([Section 2](#)) Robinson and Friday could communicate “extra-contractually” to break down the incentives that secrets tried to provide.¹² On the other hand, collusion is a problem for contracts in general. For instance, the scheme proposed by [Cremer and McLean \(1988\)](#) is not collusion-proof for similar reasons. To study collusion-proof contracts, assumptions must be made regarding coalitions' contractual ability. Assume that every coalition t maximizes a *coalitional utility function* $v_t : A \rightarrow \mathbb{R}$, quasilinear in money.¹³

Definition 4. A correlated strategy σ is *strongly enforceable* if there is a scheme $\zeta : I \times A \times S \rightarrow \mathbb{R}$ such that

$$\sum_{a_{-t}} \sigma(a) (v_t(a_{-t}, b_t) - v_t(a)) \leq \sum_{(a_{-t}, s)} \sigma(a) \sum_{i \in t} \zeta_i(a, s) (\Pr(s|a_{-t}, b_t, \rho_t) - \Pr(s|a))$$

for all $t \subset I$, $a_t \in A_t$, $(b_t, \rho_t) \in A_t \times R_t$, where $A_t = \prod_{i \in t} A_i$, $R_t = \{\rho_t : S_t \rightarrow S_t\}$, etc.

¹²The following incentive scheme deters such communication between Robinson and Friday (Friday prefers misreporting his signal to Robinson) while virtually enforcing (rest, work).

	(monitor, work)	(monitor, shirk)	(rest, work)	(rest, shirk)
g	$1/\mu, 1/\sigma$	$0, 1/\sigma$	$1/2\mu, 0$	$0, 1/2(1 - \sigma)$
b	$0, 0$	$1/(1 - \mu), 0$	$0, 1/(1 - \sigma)$	$1/2(1 - \mu), 1/2(1 - \sigma)$

¹³This assumption is standard. See for instance, [Che and Kim \(2006\)](#) and references therein. The purpose of this section is not to derive a meaningful utility for coalitions, but to use one.

Strong enforcement requires that no subset of agents can profitably deviate after sharing their information even if they can commit to sharing it non-strategically. Strong enforceability is thus especially “strong.”

We now derive the detection requirement implied by strong enforceability. For any subset of agents $t \subset I$, a *multilateral strategy* for t is any map $\sigma_t : A_t \rightarrow \Delta(A_t \times R_t)$. Intuitively, a multilateral strategy σ_t has the agents in t coordinate their deviations contingent on all recommendations to members of t . A multilateral strategy σ_t is called a *multilateral disobedience* if $\sigma_t(b_t, \rho_t | a_t) > 0$ for some (a_t, b_t, ρ_t) such that $a_t \neq b_t$. It is called *detectable* if $\Pr(s|a) \neq \Pr(s|a, \sigma_t)$ for some $a \in A$ and $s \in S$.

A *coalitional strategy* by agent i is a profile of multilateral strategies $\sigma^i = \{\sigma_t : t \ni i\}$, one for each coalition to which i may belong. It is called a *coalitional disobedience* if σ_t is a disobedience for some coalition $t \ni i$. It is called *detectable* if $\Pr(s|a) \neq \Pr(s|a, \sigma^i)$ for some $a \in A$ and $s \in S$, where

$$\Pr(s|\mu, \sigma^i) := \sum_{t \ni i} \sum_{(a, b_t, \rho_t)} \mu(a) \Pr(s|a_{-t}, b_t, \rho_t) \sigma_t(b_t, \rho_t | a_t).$$

Intuitively, a coalitional strategy for an agent i is a profile of multilateral strategies involving i . It is undetectable if for every action profile a , even if some multilateral strategy σ_t is detectable, there is another multilateral strategy $\sigma_{t'}$ with $i \in t \cap t'$ that “undoes” the change in probability from σ_t . Therefore, even if every multilateral disobedience is detectable, some coalitional disobedience may remain undetectable.

Proposition 3. *Every coalitional disobedience is detectable if and only if for any coalitional utility profile, every correlated strategy is strongly virtually enforceable.*

4.4 Genericity

We end this section by deriving conditions on the number of agents’ action-signal pairs such that every disobedience is detectable for every monitoring technology except for those in a set of Lebesgue measure zero.

Intuitively, incentives may be provided to a given agent in three ways: (a) using only others’ signals to detect his deviations (e.g., Friday), (b) using only his own reports and others’ recommendations (e.g., Robinson), and (c) using both his reports and others’ signals in conjunction. Proposition 4 below identifies conditions such that for every agent, at least one such way of detecting deviations is generic.

Proposition 4. *Every disobedience is detectable generically if for every agent i ,*

- (a) $|A_i| - 1 \leq |A_{-i}| (|S_{-i}| - 1)$ *when* $|S_i| = 1$,
- (b) $|A_i| (|S_i| - 1) \leq |A_{-i}| - 1$ *when* $|S_{-i}| = 1$, *and*
- (c) $|A_i| |S_i| \leq |A_{-i}| |S_{-i}|$ *when both* $|S_i| > 1$ *and* $|S_{-i}| > 1$.

If $|S| = 1$ then every disobedience is detectable generically only if $|A| = 1$. More interestingly, genericity holds even if $|S| = 2$, as long as agents have enough actions. Hence, a team may overcome incentive constraints generically even if only one individual can make substantive observations and these observations are just a binary bit of information. If others' action spaces are large enough and their actions have generic effect on the bit's probability, this uniquely informed individual may still be controlled by testing him with unpredictable combinations of others' actions.¹⁴

5 Literature

In this section we compare the results of this paper with the relevant literature. We begin with a discussion of the partnership problem, which motivated [Alchian and Demsetz \(1972\)](#) and others to study the role of monitoring, together with the literature on subjective evaluation. Finally, several papers with comparable detectability criteria, especially in the literature on repeated games, are discussed.

5.1 The Partnership Problem

The partnership problem was introduced by [Alchian and Demsetz \(1972\)](#) and may be described intuitively as follows. Consider two people working together in an enterprise that involves mutual effort. The efficient amount of effort would align each party's marginal effort cost with its marginal benefit, which in a competitive economy coincides with the firm's profit. However, each individual has the incentive to align his marginal effort cost with just his share of the marginal benefit, rather than the entire marginal benefit. This inevitably leads to shirking. One way to solve—or at least mitigate—this shirking problem would be for the firm to hire a monitor in order to contract directly for the workers' effort. But then who will monitor the monitor?

¹⁴I thank Roger Myerson for urging me to emphasize this point.

According to [Alchian and Demsetz \(1972, p. 778, their footnote\)](#), *[t]wo key demands are placed on an economic organization—metering input productivity and metering rewards.*¹⁵ At the heart of their “metering problem” lies the question of how to give incentives to monitors, which they answered by making the monitor residual claimant. However, this can leave the monitor with incentives to misreport input productivity if his report influences input rewards, like workers’ wages, since—given efforts—paying workers hurts him directly.¹⁶ Hence, making the monitor residual claimant, or principal, fails to provide the right incentives.

On the other hand, [Holmström \(1982, p. 325\)](#) argues that *... the principal’s role is not essentially one of monitoring ... the principal’s primary role is to break the budget-balance constraint.* He argues that if output is publicly verifiable then the principal can provide the right incentives to agents with “group penalties” that reward all agents when output is good and punish them all when it is bad. Where [Alchian and Demsetz](#) seem to overemphasize the role of monitoring in organizations, [Holmström](#) seems to underemphasize it. By assuming that output is publicly verifiable, he finds little role for monitoring,¹⁷ and as a result [Holmström \(1982, p. 339\)](#) concludes wondering: *... how should output be shared so as to provide all members of the organization (including monitors) with the best incentives to perform?*

The problem of providing incentives in organizations without publicly verifiable output has been recognized by several authors ([Prendergast, 1999](#); [MacLeod, 2003](#); [Levin, 2003](#); [Fuchs, 2007](#)) under the rubric of “subjective evaluation.” They study a firm with a “monitor” who privately observes output. To give the monitor incentives for truthful reporting, one way or another they make the monitor’s earnings independent of his report. This way, he is happy to report observed output. Presumably, his agents are rewarded for high output, so there is a wedge between the monitor’s payoff and the agents’ wages. This wedge may be burned or sold. However, the literature on subjective evaluation leaves open a basic problem: what if observing output is costly?

¹⁵*Meter means to measure and also to apportion. One can meter (measure) output and one can also meter (control) the output. We use the word to denote both; the context should indicate which.*

¹⁶A comparable argument was put forward by [Strausz \(1997\)](#) by observing that delegated monitoring dominates monitoring by a principal who cannot commit to his agent that he will verify the agent’s effort when it is only privately observed. However, [Strausz](#) assumes that monitoring signals are “hard evidence,” so a monitor cannot misreport his information. I allow for soft evidence.

¹⁷Intuitively, if output were not publicly verifiable then his group penalties would no longer provide the right incentives: monitors would always report good output to secure payment and shirk from their monitoring responsibilities to save on effort. Knowing this, workers would also shirk.

In this case, *no matter how small the cost*, rewarding a monitor independently of his report will induce him to avoid exerting effort towards reporting accurately.

This paper accommodates costly private monitoring and finds a contract that gives both workers and monitors the right incentives to perform. It also addresses the partnership problem. Even though in [Example 1](#) there was just one worker (Friday), it is easy to incorporate an additional one, call him Thursday. Following [Alchian and Demsetz \(1972\)](#), suppose that, if only one agent shirks, Robinson can tell who it was.

Both workers working is virtually enforceable with budget balance. Most of the time, the principal asks the workers to work and occasionally picks a worker at random, asks him to shirk and the other one to work. He never asks both workers to shirk. Suppose that Robinson’s report coincides with the principal’s recommendations. If he reports that both workers worked then Robinson pays both workers. If he reports that one worked and the other shirked then Robinson pays the worker but not the shirker. Now suppose that Robinson’s report differs from the recommendations. If he gets one worker’s recommendation wrong then he must pay a penalty to the worker whose recommendation he did not get wrong. If he gets both workers’ recommendations wrong then he must pay both workers a very large penalty. It is not difficult to show that this arrangement provides all agents with the right incentives to perform and the sum of payments across individuals always equals zero. (The details are available on request.) Therefore, the principal does not spend any money. Since the principal observes reports and makes recommendations at no cost, he would be happy to report the reports and recommendations truthfully even if they were not verifiable. Thus, *nobody needs to monitor the principal*. See [Rahman and Obara \(2008\)](#) for related work that characterizes enforceability with budget balance.¹⁸

Some of the literature has addressed the partnership problem from a dynamic perspective (e.g., [Radner et al., 1986](#); [Levin, 2003](#); [Fuchs, 2007](#)). This attempt adds useful specificity to the problem, because now the principal’s tools for incentive provision have a dynamic flavor, such as choosing when to dissolve the partnership. Reinterpreting continuation values as payments, we abstract from these dynamic interpretations but acknowledge that they are still implicit in this paper.

Finally, there is an important literature on market-based incentives, such as [MacLeod and Malcomson \(1998\)](#); [Prendergast \(1999\)](#); [Tadelis \(2002\)](#) and others. Although this model is not market-based, it may be incorporated into participation constraints.

¹⁸[Rahman and Obara \(2008\)](#) finds sufficient but not necessary conditions for virtual enforcement.

5.2 Detection and Enforcement

A number of papers have emphasized the duality between detection and enforcement. Some of the first references that point this out are [Abreu et al. \(1990\)](#) and [Fudenberg et al. \(1994\)](#), in the context of repeated games. Together with related literature on partnerships such as [Legros and Matsushima \(1991\)](#), [d'Aspremont and Gérard-Varet \(1998\)](#) and [Legros and Matthews \(1993\)](#), these papers restrict attention to public monitoring. Furthermore, reinterpreting continuation values as payments, none of these papers considers incentive schemes that depend on recommendations. Therefore, even though these papers characterize enforceability, they fail to enforce many outcomes that are enforceable in this paper, even in a more general context.

Some papers have considered richer contract spaces than the ones above in specific settings, such as [Kandori \(2003\)](#) and its extension to private monitoring by [Obara \(2008\)](#), [Aoyagi \(2005\)](#) and [Tomala \(2009\)](#). [Kandori \(2003\)](#), has agents play mixed strategies and report the realization of such mixtures. He considers contracts contingent on those reports and signal realizations. Mediated contracts can perform strictly better even with public monitoring, as the next example shows.

Example 5. One agent, three actions (L , M and R), and two publicly verifiable signals (g and b), with the following utility function and monitoring technology.

L	M	R	L	M	R
0	2	0	1, 0	1/2, 1/2	0, 1
Utility Payoffs			Signal Probabilities		

The mixed strategy $\sigma = \frac{1}{2}[L] + \frac{1}{2}[R]$ is enforceable with secret contracts but not with [Kandori's](#) contracts. Indeed, offering \$1 for g if asking to play L and \$1 for b if asking to play R makes σ enforceable. With [Kandori's](#) contracts, the agent is asked to play σ and then asked what he actually played before receiving any monetary rewards. The agent gains two 'utils' by playing M instead and announcing that he played L (R) if the realized signal is g (b), with the same expected monetary payoff.

It is not difficult to show that [Kandori's](#) contracts are equivalent to mediated contracts (in that they generate the same set of enforceable outcomes) if actions are *secretly* announced *before* signals are observed. In this case, the principal's recommendations need not be verifiable. (These reports must be verifiable, though.) Thus, the principal can monitor Robinson in [Example 1](#) by having Friday mix and report what he played.

Specifically, let σ be the probability that Robinson monitors and μ the probability that Friday (independently) works. When Robinson monitors, Robinson gets $\$1/\mu$ and Friday gets $\$1/\sigma$ if both agents report that Friday worked. If both agents report that Friday shirked then Robinson gets $\$1/(1 - \mu)$ and Friday gets nothing. After any other event, both agents get nothing. To virtually enforce (shirk,work), Robinson must either report what he plans to play before he plays it or what he played before he observes the signal. In this arrangement, the principal’s role is just being a budget breaker, i.e., there is no explicit need for a mediator. As long as Robinson and Friday can commit to destroy or sell value, they can write this contract by themselves.

We end this section by contrasting our work with some other relevant papers. [Aoyagi](#) finds dynamic mediated strategies that rely on “ ε -perfect” monitoring, and fail if monitoring is costly or one-sided. In a repeated game, [Tomala](#) studies a class of recursive communication equilibria and independently considers recommendation-contingent contracts with continuation values to prove a folk theorem. [Tomala](#) focuses on exact implementation and proves a version of our Minimax Lemma. However, he does not consider virtual enforcement. He defines detectability with respect to a fixed correlated strategy using unconditional probabilities over actions and signals.

Last but not least, the work of [Lehrer \(1992\)](#) is especially noteworthy. In the context of a repeated game, he characterizes the equilibrium payoffs set of a two-player game with imperfect monitoring and time-average utilities (heuristically, discount factors equal one). There are some similarities between his results and ours, but also important differences. He characterizes equilibrium payoffs as follows. A payoff profile is *sustainable* if there is a strategy profile $\mu = (\mu_1, \mu_2)$ that attains it and every μ -profitable disobedience is detectable by some behavior, not necessarily μ . This is established by having players undertake detecting behavior as the game proceeds with probability diminishing so quickly that it does not enter the time-average utility.

However, his argument relies on the use of time-average utility. In order to characterize the equilibrium payoffs set as the discount factor tends to one rather than at the limit, virtual enforcement becomes the appropriate notion. Indeed, according to [Lehrer’s](#) characterization, the profile (rest,work) should be sustainable in [Example 4](#) because every (rest,work)-profitable disobedience is detectable. However, as was pointed out, it is not virtually enforceable and furthermore the issue is more subtle. Understanding in detail how [Theorem 4](#) describes the limit of equilibrium payoffs sets as the discount factor tends to one is the object of future research.

6 Conclusion

This paper provides the following answer to [Alchian and Demsetz](#)’s question of who will monitor the monitor. First of all, it is a bad idea to make a costly private monitor the principal, so make him an agent. The principal monitors the monitor’s detectable deviations by having his workers shirk occasionally, and nobody needs to monitor the monitor’s undetectable deviations ([Theorem 3](#)). How to monitor the monitor? By asking him “trick questions” and offering him Robinson’s contract ([Example 1](#)). This contract aligns incentives by making the monitor responsible for monitoring.

[Alchian and Demsetz](#) argued that the monitor ought to be made residual claimant because only then his incentives be appropriately aligned. In a sense, they “elevated” the role of monitoring in organizations. On the other hand, in this paper the monitor is “demoted” to a security guard—low down in the ownership hierarchy. As such, the question remains: what is the economic role of residual claimant? Answering this question is the purpose of future research.

[Knight](#) (1921, Part III, Chapter IX, par. 10) aptly argues that . . . *there must come into play the diversity among men in degree of confidence in their judgment and powers and in disposition to act on their opinions, to “venture.” This fact is responsible for the most fundamental change of all in the form of organization, the system under which the confident and venturesome “assume the risk” or “insure” the doubtful and timid by guaranteeing to the latter a specified income in return for an assignment of the actual results.*

This suggests that the role of residual claimant is to screen prospective members of an organization. Indeed, according to [Knight](#) (1921, Part III, Chapter IX, par. 11) *With human nature as we know it it would be impracticable or very unusual for one man to guarantee to another a definite result of the latter’s actions without being given power to direct his work. And on the other hand the second party would not place himself under the direction of the first without such a guaranty.*

In other words, individuals claim the group’s residual in order to reassure the group that they can lead them into profitable activities, thereby separating themselves from individuals who would not be able to lead the group in the right direction. A related argument might be attributed to [Leland and Pyle](#) (1977), who argued for the signaling nature of retained equity.

A Proofs

Theorem 1. By the Alternative Theorem (see, e.g., [Rockafellar, 1970](#), Theorem 22.1), μ is not enforceable if and only if there exists a vector $\lambda \geq 0$ and an agent i such that

$$\sum_{(b_i, \rho_i)} \mu(a) \lambda_i(a_i, b_i, \rho_i) [\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a)] = 0 \quad \forall (a, s)$$

and $\Delta v_i(\mu, \lambda_i) > 0$. This vector λ exists if and only if the strategy σ_i , defined pointwise by

$$\sigma_i(b_i, \rho_i|a_i) := \begin{cases} \lambda_i(a_i, b_i, \rho_i) / \sum_{(b'_i, \rho'_i)} \lambda_i(a_i, b'_i, \rho'_i) & \text{if } \sum_{(b'_i, \rho'_i)} \lambda_i(a_i, b'_i, \rho'_i) > 0, \text{ and} \\ [(a_i, \tau_i)](b_i, \rho_i) & \text{otherwise (where } [\cdot] \text{ denotes Dirac measure),} \end{cases}$$

is μ -profitable and supp μ -undetectable. \square

Theorem 2. Let $B = \text{supp } \mu$. By the Alternative Theorem, every B -disobedience is B -detectable if and only if a scheme ξ exists such that $\xi_i(a, s) = 0$ if $a \notin B$ and

$$0 \leq \sum_{(a_{-i}, s)} \xi_i(a, s) (\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a)) \quad \forall i \in I, a_i \in B_i, b_i \in A_i, \rho_i \in R_i,$$

with a strict inequality whenever $a_i \neq b_i$, where $B_i = \{a_i \in A_i : \exists a_{-i} \in A_{-i} \text{ s.t. } a \in B\}$. Replacing $\xi_i(a, s) = \mu(a) \zeta_i(a, s)$ for any correlated strategy μ with $\text{supp } \mu = B$, this is equivalent to there being, for every v , an appropriate rescaling of ζ that satisfies (*). \square

Theorem 3. Let $B = \text{supp } \mu$. For necessity, suppose there is a B -disobedient, undetectable disobedience σ_i , so $\sigma_i(b_i, \rho_i|a_i) > 0$ for some $a_i \in B_i$, $b_i \neq a_i$ and $\rho_i \in R_i$. Letting $v_i(a_{-i}, b_i) < v_i(a)$ for every a_{-i} , clearly no correlated strategy with positive probability on a_i is virtually enforceable. Sufficiency follows by [Lemmata B.3, B.4 and B.10](#). \square

Theorem 4. For sufficiency, suppose that μ is virtually enforceable, so there is a sequence $\{\mu^m\}$ such that μ^m is enforceable for every m and $\mu^m \rightarrow \mu$. Without loss, assume that $\text{supp } \mu^m \supset \text{supp } \mu$ for all m . If $\mu^m = \mu$ for all large m then μ is enforceable and the condition of [Theorem 4](#) is fulfilled with $\tilde{\mu} = \mu$, so suppose not. If there exists m and m' such that $\mu^m = p\mu^{m'} + (1-p)\mu$ then incentive compatibility with respect to m yields that $\sum_{a_{-i}} \mu^m(a) \Delta v_i(a, \sigma_i) \leq \sum_{a_{-i}} \mu^m(a) \zeta_i^m(a) \cdot \Delta \Pr(a, \sigma_i) \leq \sum_{a_{-i}} \mu^m(a) \bar{z} \|\Delta \Pr(a, \sigma_i)\|$ for every σ_i , where $\bar{z} = \max_{(i, a, s)} |\zeta_i^m(a, s)| + 1$ and ζ^m enforces μ^m for each m . For large m' , $\mu^{m'}$ is sufficiently close to μ that if σ_i is μ -profitable then $\sum_{a_{-i}} \mu^{m'}(a) \Delta v_i(a, \sigma_i) > 0$, so σ_i is detectable. Therefore, $\sum_{a_{-i}} \mu^m(a) \Delta v_i(a, \sigma_i) < \sum_{a_{-i}} \mu^m(a) \bar{z} \|\Delta \Pr(a, \sigma_i)\|$.

If there does not exist m and m_1 such that $\mu^m = p\mu^{m_1} + (1-p)\mu$ then there exists μ^{m_2} such that its distance from μ is less than the positive minimum distance between μ and the affine hull of $\{\mu^m, \mu^{m_1}\}$. Therefore, the lines generated by μ^m and μ^{m_1} and μ^{m_1} and μ^{m_2} are not

collinear. Proceeding inductively, pick $C = \{\mu^{m_1}, \dots, \mu^{m_{|A|}}\}$ such that its affine space is full-dimensional in $\Delta(A)$. Since we are assuming that μ is not enforceable, it lies outside $\text{conv } C$. Let $\hat{\mu} = \sum_k \mu^{m_k} / |A|$ and $B_\varepsilon(\hat{\mu})$ be the open ε -ball around $\hat{\mu}$ for some $\varepsilon > 0$. By construction, $B_\varepsilon(\hat{\mu}) \subset \text{conv } C$ for $\varepsilon > 0$ sufficiently small, so there exists $\hat{\mu}' \in B_\varepsilon(\hat{\mu})$ such that $p\hat{\mu} + (1-p)\mu = \hat{\mu}'$ for some p such that $0 < p < 1$. Now we can apply the argument from the previous paragraph, so the condition of [Theorem 4](#) holds.

For necessity, if μ is not virtually enforceable then $1 \geq V_\mu(z) \geq C > 0$ for every z , where V_μ is defined in [Lemma B.3](#). Let (λ^z, μ^z) solve $V_\mu(z)$ for every z . Given $\mu \in \Delta(A)$,

$$C \leq V_\mu(z) \leq 1 + \sum_{(i,a)} \Delta v_i(\mu, \lambda_i^z) - z \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\|.$$

If the condition of [Theorem 4](#) holds then $\sum_{(i,a)} \Delta v_i(\mu, \lambda_i^z) < \bar{z} \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\|$ and $\sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\| > 0$, since there must exist i such that λ_i^μ is μ -profitable. Hence, $C \leq 1 + (\bar{z} - z) \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\|$, i.e., $z - \bar{z} \leq (1 - c) / \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\|$. This inequality must hold for every z , therefore $\sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i^z)\| \rightarrow 0$ as $z \rightarrow \infty$. But this contradicts [Lemma B.11](#), since $\sum_i \Delta v_i(\mu, \lambda_i) \geq C$, completing the proof. \square

Proposition 2. Fix any $\mu \in \Delta(A)$. By the Alternative Theorem, every strategy that is undetectable at μ is μ -unprofitable if and only if a scheme $\zeta : I \times S \rightarrow \mathbb{R}$ exists such that

$$\sum_{a_{-i}} \mu(a) [v_i(a_{-i}, b_i) - v_i(a)] \leq \sum_{(a_{-i}, s)} \mu(a) \zeta_i(s) [\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a)] \quad \forall (i, a_i, b_i, \rho_i).$$

The result now follows. This proof follows closely that of [Theorem 1](#). \square

Proposition 3. This proof is similar to the previous ones, applying the Alternative Theorem, except that now the incentive constraints include multilateral deviations. \square

Proposition 4. By [Lemma B.1](#), DUD is implied by *conic independence*

$$\forall (i, a_i, s_i), \quad \Pr(a_i, s_i) \notin \text{cone}\{\Pr(b_i, t_i) : (b_i, t_i) \neq (a_i, s_i)\}.$$

This is in turn implied by *linear independence*, or full row rank, for all i , of the matrix with $|A_i| |S_i|$ rows, $|A_{-i}| |S_{-i}|$ columns and entries $\Pr(a_i, s_i)(a_{-i}, s_{-i}) = \Pr(s|a)$. Since the set of full rank matrices is generic, this full row rank is generic if $|A_i| |S_i| \leq |A_{-i}| |S_{-i}|$ if $|S_i| > 1$ and $|S_{-i}| > 1$. If $|S_i| = 1$, adding with respect to s_{-i} for each a_{-i} yields column vectors equal to $(1, \dots, 1) \in \mathbb{R}^{A_i}$. This leaves $|A_{-i}| - 1$ linearly dependent columns. Eliminating them, genericity requires that for every i ,

$$|A_i| = |A_i| |S_i| \leq |A_{-i}| |S_{-i}| - (|A_{-i}| - 1) = |A_{-i}| \times (|S_{-i}| - 1) + 1.$$

Similarly, there are $|A_i| - 1$ redundant row vectors when $|S_{-i}| = 1$. Since the intersection of finitely many generic sets is generic, DUD is generic if all these conditions hold. \square

B Lemmata

Lemma B.1. *Every disobedience is detectable if*

$$\forall(i, a_i, s_i), \quad \Pr(a_i, s_i) \notin \text{cone}\{\Pr(b_i, t_i) : (b_i, t_i) \neq (a_i, s_i)\},$$

where cone stands for the set of positive linear combinations of $\{\Pr(b_i, t_i) : (b_i, t_i) \neq (a_i, s_i)\}$.

Proof. If DUD fails then there exists σ_i such that $\sigma_i(b_i, \rho_i|a_i) > 0$ for some $a_i \neq b_i$ and

$$\begin{aligned} \forall(a, s), \quad \Pr(s|a) &= \sum_{(b_i, \rho_i)} \sum_{t_i \in \rho_i^{-1}(s_i)} \sigma_i(b_i, \rho_i|a_i) \Pr(s_{-i}, t_i|a_{-i}, b_i) \\ &= \sum_{(b_i, t_i)} \sum_{\{\rho_i: \rho_i(t_i)=s_i\}} \sigma_i(b_i, \rho_i|a_i) \Pr(s_{-i}, t_i|a_{-i}, b_i). \end{aligned}$$

Write $\lambda_i(a_i, s_i, b_i, t_i) := \sum_{\{\rho_i: \rho_i(t_i)=s_i\}} \sigma_i(b_i, \rho_i|a_i)$. By construction, $\lambda_i(a_i, s_i, b_i, t_i) \geq 0$ is strictly positive for some $a_i \neq b_i$ and satisfies

$$\forall(i, a, s), \quad \Pr(s|a) = \sum_{(b_i, t_i)} \lambda_i(a_i, s_i, b_i, t_i) \Pr(s_{-i}, t_i|a_{-i}, b_i).$$

Without loss, $\lambda_i(a_i, s_i, a_i, s_i) = 0$ for some (a_i, s_i) . Indeed, if $\lambda_i(a_i, s_i, a_i, s_i) = 1$ for all (a_i, s_i) , then the equation above is violated because σ_i is disobedient by hypothesis and probabilities are non-negative. If $\lambda_i(a_i, s_i, a_i, s_i) \neq 1$ then subtract $\lambda_i(a_i, s_i, a_i, s_i) \Pr(s|a)$ from both sides of the equation and divide by $1 - \lambda_i(a_i, s_i, a_i, s_i)$. Therefore, $\Pr(a_i, s_i) \in \text{cone}\{\Pr(b_i, t_i) : (b_i, t_i) \neq (a_i, s_i)\}$ for some (a_i, s_i) . \square

Let $\mathcal{D}_i = \Delta(A_i \times R_i)^{A_i}$ be the space of strategies σ_i for a agent i and $\mathcal{D} = \prod_i \mathcal{D}_i$ the set of strategy profiles $\sigma = (\sigma_1, \dots, \sigma_n)$. Call μ *enforceable within* some vector $z \in \mathbb{R}_+^I$ if there is a scheme ξ that satisfies $(*)$ and $-\mu(a)z_i \leq \xi_i(a, s) \leq \mu(a)z_i$ for all (i, a, s) . Next, we provide a lower bound on z so that μ is enforceable within z .

Lemma B.2. (i) *A correlated strategy μ is enforceable within $z \in \mathbb{R}_+^I$ if and only if*

$$V_\mu(z) := \max_{\sigma \in \mathcal{D}} \sum_{i \in I} \Delta v_i(\mu, \sigma_i) - \sum_{(i, a)} z_i \mu(a) \|\Delta \Pr(a, \sigma_i)\| = 0.$$

(ii) *If μ is enforceable then $V_\mu(z) = 0$ for some $z \in \mathbb{R}_+^I$. If not then $\sup_z V_\mu(z) > 0$.*

(iii) *A correlated strategy μ is enforceable if and only if $\bar{z}_i < +\infty$ for every agent i , where*

$$\bar{z}_i := \sup_{\sigma_i \in \mathcal{F}_i} \frac{\max\{\Delta v_i(\mu, \sigma_i), 0\}}{\sum_a \mu(a) \|\Delta \Pr(a, \sigma_i)\|} \quad \text{if } \mathcal{F}_i := \{\sigma_i : \sum_a \mu(a) \|\Delta \Pr(a, \sigma_i)\| > 0\} \neq \emptyset$$

and, whenever $\mathcal{F}_i = \emptyset$, $\bar{z}_i := +\infty$ exactly when $\max_{\sigma_i} \Delta v_i(\mu, \sigma_i) > 0$.¹⁹

(iv) *If $\bar{z}_i < +\infty$ for every i then $V_\mu(z) = 0$ if and only if $z_i \geq \bar{z}_i$ for all i .*

¹⁹Intuitively, \mathcal{F}_i is the set of all supp μ -detectable deviation plans available to agent i .

Proof. Consider the family of linear programs below indexed by $z \in [0, \infty)^I$.

$$\begin{aligned} \max_{\varepsilon \geq 0, \xi} \quad & - \sum_{(i, a_i)} \varepsilon_i(a_i) \quad \text{s.t.} \quad \forall(i, a, s), \quad -\mu(a)z_i \leq \xi_i(a, s) \leq \mu(a)z_i, \\ & \forall(i, a_i, b_i, \rho_i), \quad \sum_{a_{-i}} \mu(a) \Delta v_i(a, b_i) - \sum_{a_{-i}} \xi_i(a) \cdot \Delta \Pr(a, b_i, \rho_i) \leq \varepsilon_i(a_i), \end{aligned}$$

where $\Delta v_i(a, b_i) := v_i(a_{-i}, b_i) - v_i(a)$ and $\Delta \Pr(a, b_i, \rho_i) := \Pr(a_{-i}, b_i, \rho_i) - \Pr(a)$. Given $z \geq 0$, the primal problem above looks for a scheme ξ adapted to μ (i.e., such that $\xi_i(a, s) = 0$ whenever $\mu(a) = 0$) that minimizes the burden $\varepsilon_i(a_i)$ of relaxing incentive constraints. By construction, μ is enforceable with transfers bounded by z if and only if there is a feasible ξ with $\varepsilon_i(a_i) = 0$ for all (i, a_i) , i.e., the value of the problem is zero. Since μ is assumed enforceable, such z exists. The dual of this problem is:

$$\begin{aligned} \min_{\sigma, \beta \geq 0} \quad & \sum_{(i, a)} \mu(a) [z_i \sum_{s \in S} \mu(a) (\beta_i^+(a, s) + \beta_i^-(a, s)) - \Delta v_i(a, \sigma_i)] \quad \text{s.t.} \\ & \forall(i, a_i), \quad \sum_{(b_i, \rho_i)} \sigma_i(b_i, \rho_i | a_i) \leq 1, \\ & \forall i \in I, a \in \text{supp } \mu, s \in S, \quad \Delta \Pr(s|a, \sigma_i) = \beta_i^+(a, s) - \beta_i^-(a, s). \end{aligned}$$

Since $\beta_i^\pm(a, s) \geq 0$, it follows easily that $\beta_i^+(a, s) = \max\{\Delta \Pr(s|a, \sigma_i), 0\}$ and $\beta_i^-(a, s) = \min\{\Delta \Pr(s|a, \sigma_i), 0\}$. Hence, $\beta_i^+(a, s) + \beta_i^-(a, s) = |\Delta \Pr(s|a, \sigma_i)|$. Since $\|\Delta \Pr(a, \sigma_i)\| = \sum_s |\Delta \Pr(s|a, \sigma_i)|$, the dual is now equivalent to

$$V_\mu(z) = \max_{\sigma \geq 0} \sum_{(i, a)} \mu(a) (\Delta v_i(a, \sigma_i) - z \|\Delta \Pr(a, \sigma_i)\|) \quad \text{s.t.} \quad \forall(i, a_i), \quad \sum_{(b_i, \rho_i)} \sigma_i(b_i, \rho_i | a_i) \leq 1.$$

Adding mass to $\sigma_i(a_i, \tau_i | a_i)$ if necessary, without loss σ_i is a deviation plan, proving (i).

To prove (ii), the first sentence is obvious. The second follows by [Theorem 1](#): if μ is not enforceable then a μ -profitable, supp μ -undetectable plan σ_i exists, so $V_\mu(z) > 0$ for all z .

For (iii), if μ is not enforceable then there is a μ -profitable, supp μ -undetectable deviation plan σ_i^* . Approaching σ_i^* from \mathcal{F}_i (e.g., with mixtures of σ_i^* and a fixed plan in \mathcal{F}_i), the denominator defining \bar{z}_i tends to zero whilst the numerator tends to a positive amount, so \bar{z}_i is unbounded. Conversely, suppose μ is enforceable. If the sup defining \bar{z}_i is attained, we are done. If not, it is approximated by a sequence of supp μ -detectable deviation plans that converge to a supp μ -undetectable one. Since μ is enforceable, the limit is unprofitable. Let

$$F_i^\mu(\delta) := \min_{\lambda_i \geq 0} \sum_{a \in A} \mu(a) \|\Delta \Pr(a, \lambda_i)\| \quad \text{s.t.} \quad \Delta v_i(\mu, \lambda_i) \geq \delta.$$

Since every μ -profitable deviation plan is detectable by [Theorem 1](#), it follows that $F_i^\mu(\delta) > 0$ for all $\delta > 0$, and $\bar{z}_i = (\lim_{\delta \downarrow 0} F_i^\mu(\delta)/\delta)^{-1}$. Hence, it suffices to show $\lim_{\delta \downarrow 0} F_i^\mu(\delta)/\delta > 0$.

To this end, by adding variables like β above, the dual problem for F_i^μ is equivalent to:

$$\begin{aligned} F_i^\mu(\delta) &= \max_{\varepsilon \geq 0, x_i} \varepsilon \delta \quad \text{s.t.} \quad \forall(a, s), \quad -1 \leq x_i(a, s) \leq 1, \\ \forall(a_i, b_i, \rho_i), \quad &\sum_{a-i} \mu(a) (\varepsilon \Delta v_i(a, b_i) - x_i(a) \cdot \Delta \Pr(a, b_i, \rho_i)) \leq 0. \end{aligned}$$

Since μ is enforceable, there is a feasible solution to this dual (ε, x_i) with $\varepsilon > 0$. Hence, $F_i^\mu(\delta) \geq \varepsilon \delta$ for all $\delta > 0$, therefore $\lim_{\delta \downarrow 0} F_i^\mu(\delta)/\delta > 0$, as claimed.

To prove (iv), suppose that $\bar{z}_i < \infty$ for all i . We claim $V_\mu(\bar{z}) = 0$. Indeed, given $\sigma_i^* \in \mathcal{F}_i$ for all i , substituting the definition of \bar{z}_i into the objective of the minimization in (i),

$$\sum_{i \in I} \Delta v_i(\mu, \sigma_i^*) - \sum_{(i, a)} \mu(a) \sup_{\sigma_i \in \mathcal{F}_i} \left\{ \frac{\max\{\Delta v_i(\mu, \sigma_i), 0\}}{\sum_a \mu(a) \|\Delta \Pr(a, \sigma_i)\|} \right\} \|\Delta \Pr(a, \sigma_i^*)\| \leq 0.$$

If $\sigma_i^* \notin \mathcal{F}_i$ then, since μ is enforceable, every supp μ -undetectable deviation plan is unprofitable, so again the objective is non-positive, hence $V_\mu(\bar{z}) = 0$. Clearly, V_μ decreases with z , so it remains to show that $V_\mu(\bar{z}) > 0$ if $z_i < \bar{z}_i$ for some i . But by definition of \bar{z} , there is a deviation plan σ_i^* with $\Delta v_i(\mu, \sigma_i^*) / \sum_a \mu(a) \|\Delta \Pr(a, \sigma_i^*)\| > z_i$, so $V_\mu(z) > 0$. \square

Lemma B.3. *Consider the following linear program.*

$$\begin{aligned} V_\mu(z) &:= \min_{\eta \geq 0, p, \xi} p \quad \text{s.t.} \quad \sum_{a \in A} \eta(a) = p, \\ \forall(i, a, s), \quad &-(\eta(a) + (1-p)\mu(a))z \leq \xi_i(a, s) \leq (\eta(a) + (1-p)\mu(a))z, \\ \forall(i, a_i, b_i, \rho_i), \quad &\sum_{a-i} (\eta(a) + (1-p)\mu(a)) \Delta v_i(a, b_i) \leq \sum_{a-i} \xi_i(a) \cdot \Delta \Pr(a, b_i, \rho_i). \end{aligned}$$

The correlated strategy μ is approximately enforceable if and only if $V_\mu(z) \rightarrow 0$ as $z \rightarrow \infty$. The dual of the above linear program is given by the following problem:

$$\begin{aligned} V_\mu(z) &= \max_{\lambda \geq 0, \kappa} \sum_{i \in I} \Delta v_i(\mu, \lambda_i) - z \sum_{(i, a)} \mu(a) \|\Delta \Pr(a, \lambda_i)\| \quad \text{s.t.} \\ \forall a \in A, \quad &\kappa \leq \sum_{i \in I} \Delta v_i(a, \lambda_i) - z \sum_{i \in I} \|\Delta \Pr(a, \lambda_i)\|, \\ &\sum_{i \in I} \Delta v_i(\mu, \lambda_i) - z \sum_{(i, a)} \mu(a) \|\Delta \Pr(a, \lambda_i)\| = 1 + \kappa. \end{aligned}$$

Proof. The first family of primal constraints require ξ to be adapted to $\eta + (1-p)\mu$, so for any z , (η, p, ξ) solves the primal if and only if $\eta + (1-p)\mu$ is exactly enforceable with ξ . (Since correlated equilibrium exists, the primal constraint set is clearly nonempty, and for finite z it is also clearly bounded). The first statement now follows. The second statement follows by a lengthy but standard manipulation of the primal to obtain the above dual. \square

Lemma B.4. Consider the following family of linear programs indexed by $\varepsilon > 0$ and $z \geq 0$.

$$F_\mu^\varepsilon(z) := \max_{\lambda \geq 0} \min_{\eta \in \Delta(A)} \sum_{i \in I} \Delta v_i(\eta, \lambda_i) - z \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| \quad \text{s.t.} \\ \sum_{i \in I} \Delta v_i(\mu, \lambda_i) - z \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i)\| \geq \varepsilon.$$

$F_\mu^\varepsilon(z) \rightarrow -\infty$ as $z \rightarrow \infty$ for some $\varepsilon > 0$ if and only if μ is approximately enforceable.

Proof. The dual of the problem defining $F_\mu^\varepsilon(z)$ is

$$F_\mu^\varepsilon(z) = \min_{\delta, \eta \geq 0, x} -\delta \varepsilon \quad \text{s.t.} \quad \sum_{a \in A} \eta(a) = 1, \\ \forall(i, a, s), \quad -(\eta(a) + \delta \mu(a))z \leq x_i(a, s) \leq (\eta(a) + \delta \mu(a))z, \\ \forall(i, a_i, b_i, \rho_i), \quad \sum_{a-i} (\eta(a) + \delta \mu(a)) \Delta v_i(a, b_i) \leq \sum_{a-i} x_i(a) \cdot \Delta \Pr(a, b_i, \rho_i).$$

Since clearly $\varepsilon > 0$ does not affect the dual feasible set, if $F_\mu^\varepsilon(z) \rightarrow -\infty$ for some $\varepsilon > 0$ then there exists $z \geq 0$ such that $\delta > 0$ is feasible, and $\delta \rightarrow \infty$ as $z \rightarrow \infty$. Therefore, $F_\mu^\varepsilon(z) \rightarrow -\infty$ for every $\varepsilon > 0$. If $V_\mu(z) = 0$ for some z we are done by monotonicity of V_μ . Otherwise, suppose that $V_\mu(z) > 0$ for all $z > 0$. Let (λ, κ) be an optimal dual solution for $V_\mu(z)$ in Lemma B.3. By optimality, $\kappa = \min_{\eta \in \Delta(A)} \sum_i \Delta v_i(\eta, \lambda_i) - z \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\|$. Therefore, by the second dual constraint in $V_\mu(z)$ of Lemma B.3,

$$V_\mu(z) = 1 + \kappa = 1 + F_\mu^{V_\mu(z)}(z) = 1 - \delta V_\mu(z),$$

where δ is an optimal solution to the dual with $\varepsilon = V_\mu(z)$. Rearranging, $V_\mu(z) = 1/(1 + \delta)$. Finally, $F_\mu^\varepsilon(z) \rightarrow -\infty$ as $z \rightarrow \infty$ if and only if $\delta \rightarrow \infty$, if and only if $V_\mu(z) \rightarrow 0$. \square

Lemma B.5. Fix any $\varepsilon > 0$ and let $B = \text{supp } \mu$. If every B -disobedience is detectable then for every $C \leq 0$ there exists $z \geq 0$ such that $G_\mu(z) \leq C$, where

$$\Delta v_i(a_i)^* := \max_{(a-i, b_i)} \{\Delta v_i(a, b_i)\}, \quad \Delta v_i(a_i, \lambda_i)^* := \Delta v_i(a_i)^* \sum_{(a_i, b_i \neq a_i, \rho_i)} \lambda_i(a_i, b_i, \rho_i), \text{ and} \\ G_\mu(z) := \max_{\lambda \geq 0} \sum_{(i,a)} \|\Delta v_i(a_i, \lambda_i)\| - z \sum_{(i,a)} \|\Delta \Pr(a, \lambda_i)\| \quad \text{s.t.} \\ \forall i \in I, a_i \notin B_i, \lambda_i(a_i) = 0, \quad \text{and} \quad \sum_{i \in I} \Delta v_i(\mu, \lambda_i) - z \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i)\| \geq \varepsilon.$$

Proof. The dual of this problem is given by

$$G_\mu(z) = \min_{\delta \geq 0, x} -\delta \varepsilon \quad \text{s.t.} \\ \forall(i, a, s), \quad -(1 + \delta \mu(a))z \leq x_i(a, s) \leq (1 + \delta \mu(a))z, \\ \forall(i, a_i \in B_i, b_i, \rho_i), \quad \sum_{a-i} \delta \mu(a) \Delta v_i(a, b_i) + \mathbf{1}_{\{a_i \neq b_i\}} \Delta v_i(a_i)^* \leq \sum_{a-i} x_i(a) \cdot \Delta \Pr(a, b_i, \rho_i),$$

where $\mathbf{1}_{\{b_i \neq a_i\}} = 1$ if $b_i \neq a_i$ and 0 otherwise. This problem looks almost exactly like the dual for $F_\mu^\varepsilon(z)$ except that the incentive constraints are only indexed by $a_i \in B_i$. Now, every B -disobedience is detectable if and only if there is an incentive scheme x such that

$$0 \leq \sum_{a_{-i}} x_i(a) \cdot \Delta \Pr(a, b_i, \rho_i) \quad \forall (i, a_i, b_i, \rho_i),$$

with a strict inequality whenever $a_i \in B_i$ and $a_i \neq b_i$. Hence, by scaling x appropriately, there is a feasible dual solution with $\delta > 0$, so $G_\mu(z) < 0$. Moreover, for any $\delta > 0$, it follows that an x exists with $\sum_{a_{-i}} \delta \mu(a) \Delta v_i(a, b_i) + \mathbf{1}_{\{b_i \neq a_i\}} \Delta v_i(a_i)^* \leq \sum_{a_{-i}} x_i(a) \cdot \Delta \Pr(a, b_i, \rho_i)$ on all $(i, a_i \in B_i, b_i, \rho_i)$, so there exists z to make such δ feasible. In particular, $\delta \geq C/\varepsilon$ is feasible for some z , as required. \square

Lemma B.6. *If every B -disobedience is detectable then there exists a finite $z \geq 0$ such that*

$$\forall i \in I, a_i \in B_i, \lambda_i \geq 0, \quad \sum_{a_{-i}} \Delta v_i(a_i, \lambda_i)^* - z \|\Delta \Pr(a, \lambda_i)\| \leq 0.$$

Proof. Given $i, a_i \in B_i$, plug $\mu(a) = 1/|A_{-i}|$ for all a_{-i} in the proof of Lemma B.2 (iii). \square

Call λ *extremely detectable* if $\lambda_i(a_i)$ cannot be written as a positive linear combination involving undetectable deviations (possibly mixed) for every (i, a_i) . Let \mathcal{E} denote the set of all such extremely detectable λ .

Lemma B.7. *The set $\mathcal{D}^e = \{\sigma \in \mathcal{E} : \forall (i, a_i), \sum_{(b_i, \rho_i)} \sigma_i(a_i, b_i, \rho_i) = 1\}$ is compact.*

Proof. \mathcal{D}^e is clearly a bounded subset of Euclidean space, so it remains to show that it is closed. Consider a sequence $\{\sigma^m\} \subset \mathcal{D}^e$ such that $\sigma^m \rightarrow \sigma^*$. For any $\sigma \in \mathcal{D}$, let

$$p^*(\sigma) := \max_{0 \leq p \leq 1, \sigma^i \in \mathcal{D}} \{p : \sigma^0 \text{ is undetectable, } p\sigma^0 + (1-p)\sigma^1 = \sigma\}.$$

This is a well-defined linear program with a compact constraint set and finite values, so p^* is continuous in σ . By assumption, $p^*(\sigma^m) = 0$ for all m , so $p^*(\sigma^*) = 0$, hence $\sigma^* \in \mathcal{D}^e$. \square

Lemma B.8. *Let \mathcal{D}^e be the set of extremely detectable deviation plans.*

$$\gamma := \min_{\sigma^e \in \mathcal{D}^e} \sum_{(i, a)} \|\Delta \Pr(a, \sigma_i^e)\| > 0.$$

Proof. If $\mathcal{D}^e = \emptyset$ then $\gamma = +\infty$. If not, \mathcal{D}^e is compact by Lemma B.7, so there is no sequence $\{\sigma_i^{e,m}\} \subset \mathcal{D}^e$ with $\|\Delta \Pr(a, \sigma_i^{e,m})\| \rightarrow 0$ for all (i, a) as $m \rightarrow \infty$, hence $\gamma > 0$. \square

Lemma B.9. *Let $\mathcal{D}_i^e = \text{proj}_i \mathcal{D}^e$. There exists a finite $z \geq 0$ such that*

$$\forall i \in I, a_i \notin B_i, \sigma_i^e \in \mathcal{D}_i^e, \quad \sum_{a_{-i}} \Delta v_i(a_i, \sigma_i^e)^* - z \|\Delta \Pr(a, \sigma_i^e)\| \leq 0.$$

Proof. Let $\|\Delta v\| = \max_{(i,a,b_i)} |\Delta v_i(a, b_i)|$. If $z \geq \|\Delta v\| / \gamma$, with γ as in Lemma B.8, then

$$\forall(i, a_i), \quad \sum_{a_{-i}} \Delta v_i(a_i, \sigma_i^e)^* - z \|\Delta \Pr(a, \sigma_i^e)\| \leq \|\Delta v\| - z \sum_{a_{-i}} \|\Delta \Pr(a, \sigma_i^e)\| \leq \|\Delta v\| - \frac{\|\Delta v\|}{\gamma} \gamma.$$

The right-hand side clearly equals zero, which establishes the claim. \square

Lemma B.10. *Fix any $\varepsilon > 0$. If every B-disobedience is detectable then for every $C \leq 0$ there exists $z \geq 0$ such that for every $\lambda \geq 0$ with*

$$\sum_{i \in I} \Delta v_i(\mu, \lambda_i) - z \sum_{(i,a)} \mu(a) \|\Delta \Pr(a, \lambda_i)\| \geq \varepsilon,$$

there exists $\eta \in \Delta(A)$ such that

$$W(\eta, \lambda) := \sum_{i \in I} \Delta v_i(\eta, \lambda_i) - z \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| \leq C.$$

Proof. Rewrite $W(\eta, \lambda)$ by splitting it into three parts, $W_d(\eta, \lambda)$, $W_e(\eta, \lambda)$ and $W_u(\eta, \lambda)$:

$$\begin{aligned} W_d(\eta, \lambda) &= \sum_{i \in I} \sum_{a_i \in B_i} \sum_{a_{-i}} \eta(a) (\Delta v_i(a, \lambda_i) - z \|\Delta \Pr(a, \lambda_i)\|) \\ W_e(\eta, \lambda) &= \sum_{i \in I} \sum_{a_i \notin B_i} \sum_{a_{-i}} \eta(a) (\Delta v_i(a, \lambda_i^e) - z \|\Delta \Pr(a, \lambda_i^e)\|), \\ W_u(\eta, \lambda) &= \sum_{i \in I} \sum_{a_i \notin B_i} \sum_{a_{-i}} \eta(a) (\Delta v_i(a, \lambda_i^u) - z \|\Delta \Pr(a, \lambda_i^u)\|), \end{aligned}$$

and $\lambda = \lambda^e + \lambda^u$ with λ^e extremely detectable, λ^u undetectable. Since λ^u is undetectable,

$$W_u(\eta, \lambda) = \sum_{i \in I} \sum_{a_i \notin B_i} \sum_{a_{-i}} \eta(a) \Delta v_i(a, \lambda_i^u)$$

Let $\eta^0(a) = 1/|A|$ for every a . By Lemma B.5, there exists z with $W_d(\eta^0, \lambda) \leq C$ for every λ , and by Lemma B.9 there exists z with $W_e(\eta^0, \lambda) \leq 0$ for every λ . Therefore, if $W_u(\eta^0, \lambda) \leq 0$ we are done. Otherwise, for every i and $a_i, b_i \in A_i$, let $\eta_i^0(a_i) = 1/|A_i|$ and

$$\eta_i^1(b_i) := \sum_{(a_i, \rho_i)} \frac{\lambda_i^u(a_i, b_i, \rho_i)}{\sum_{(b'_i, \rho'_i)} \lambda_i^u(a_i, b'_i, \rho'_i)} \eta_i^0(a_i)$$

Iterate this rule to obtain a sequence $\{\eta_i^m\}$ with limit $\eta_i^\infty \in \Delta(A_i)$. By construction, η_i^∞ is a λ_i^u -stationary distribution (Nau and McCardle, 1990; Myerson, 1997). Therefore, given any a_{-i} , the deviation gains for every agent equal zero, i.e.,

$$\sum_{(a_i, b_i, \rho_i)} \eta_i^\infty(a_i) \lambda_i^u(a_i, b_i, \rho_i) (v_i(a_{-i}, b_i) - v_i(a)) = 0.$$

Let $\eta^m(a) := \prod_i \eta_i^m(a_i)$ for all m . By construction, $W_u(\eta^\infty, \lambda^u) = 0$. We will show that $W_d(\eta^\infty, \lambda) \leq C$ and $W_e(\eta^\infty, \lambda) \leq 0$. To see this, notice firstly that, since λ_i^u is undetectable, for any other agent $j \neq i$, any $\lambda_j \geq 0$ and every action profile $a \in A$,

$$\|\Delta \Pr(a, \lambda_j)\| = \|\Delta \Pr(a, \lambda_j^u, \lambda_j)\| \leq \|\Delta \Pr(a, \hat{\lambda}_j^u, \lambda_j)\|,$$

where $\hat{\lambda}_i^u(a_i, b_i, \tau_i) = \sum_{\rho_i} \lambda_i^u(a_i, b_i, \rho_i)$ and $\hat{\lambda}_i^u(a_i, b_i, \rho_i) = 0$ for all $\rho_i \neq \tau_i$,

$$\Delta \Pr(a, \lambda_i^u, \lambda_j) = \sum_{(b_j, \rho_j)} \lambda_j(a_j, b_j, \rho_j) \sum_{(b_i, \rho_i)} \lambda_i^u(a_i, b_i, \rho_i) (\Pr(a, b_i, \rho_i, b_j, \rho_j) - \Pr(a, b_i, \rho_i)),$$

and $\Pr(s|a, b_i, \rho_i, b_j, \rho_j) = \sum_{t_j \in \rho_j^{-1}(s_j)} \Pr(s_{-j}, t_j|a_{-j}, b_j, b_i, \rho_i)$. Secondly, notice that

$$\begin{aligned} \forall i \in I, a_i \in B_i, \quad & \sum_{a_{-i}} \eta^m(a) (\Delta v_i(a, \lambda_i) - z \|\Delta \Pr(a, \lambda_i)\|) \leq \\ & \eta_i^m(a_i) \sum_{a_{-i}} \eta_{-i}^m(a_{-i}) (\Delta v_i(a_i, \lambda_i)^* - z \|\Delta \Pr(a, \lambda_i)\|) \leq \\ & \eta_i^m(a_i) \sum_{a_{-i}} \eta_{-i}^0(a_{-i}) (\Delta v_i(a_i, \lambda_i)^* - z \|\Delta \Pr(a, \lambda_i)\|) \leq \\ & \sum_{a_{-i}} \eta^0(a) (\Delta v_i(a_i, \lambda_i)^* - z \|\Delta \Pr(a, \lambda_i)\|). \end{aligned}$$

Indeed, the first inequality is obvious. The second one follows by repeated application of the previously derived inequality $\|\Delta \Pr(a, \lambda_i)\| \leq \|\Delta \Pr(a, \hat{\lambda}_j^u, \lambda_i)\|$ for each agent $j \neq i$ separately m times. The third inequality follows because (i) $\eta_i^m(a_i) \geq \eta_i^0(a_i)$ for all m and $a_i \in B_i$, since B_i is a $\hat{\lambda}_i^u$ -absorbing set, and (ii) $\sum_{a_{-i}} \Delta v_i(a_i, \lambda_i)^* - z \|\Delta \Pr(a, \lambda_i)\| \leq 0$ for every (i, a_i) by [Lemma B.6](#). Therefore, $W_d(\eta^\infty, \lambda) \leq W_d(\eta^m, \lambda) \leq W_d(\eta^0, \lambda) \leq C$. Thirdly,

$$\begin{aligned} \forall i \in I, a_i \notin B_i, \quad & \sum_{a_{-i}} \eta_{-i}^m(a_{-i}) (\Delta v_i(a, \lambda_i^e) - z \|\Delta \Pr(a, \lambda_i^e)\|) \leq \\ & \sum_{a_{-i}} \eta_{-i}^m(a_{-i}) (\Delta v_i(a_i, \lambda_i^e)^* - z \|\Delta \Pr(a, \lambda_i^e)\|) \leq \\ & \sum_{a_{-i}} \eta_{-i}^0(a_{-i}) (\Delta v_i(a_i, \lambda_i^e)^* - z \|\Delta \Pr(a, \lambda_i^e)\|) \leq 0. \end{aligned}$$

The first inequality is again obvious, the second inequality follows by repeated application of $\|\Delta \Pr(a, \lambda_i)\| \leq \|\Delta \Pr(a, \hat{\lambda}_j^u, \lambda_i)\|$, and the third one follows from [Lemma B.9](#). Hence, $W_e(\eta^m, \lambda) \leq 0$ for every m , therefore $W_e(\eta^\infty, \lambda) \leq 0$. This completes the proof. (This proof extends [Nau and McCardle \(1990\)](#) and [Myerson \(1997\)](#) by including transfers.) \square

Lemma B.11. *The conditions of [Theorem 4](#) imply that for every $\varepsilon > 0$ there exists $\delta > 0$ such that $\sum_i \Delta v_i(\mu, \lambda_i) \geq \varepsilon$ implies that $\sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| \geq \delta$ for some $\eta \in \Delta(A)$ with $\sum_i \Delta v_i(\eta, \lambda_i) \leq \bar{z} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\|$.*

Proof. Otherwise, there exists $\varepsilon > 0$ such that for every $\delta > 0$ some λ^δ exists with $\sum_i \Delta v_i(\mu, \lambda_i^\delta) \geq \varepsilon$ but $\sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| < \delta$ whenever $\eta \in \Delta(A)$ satisfies the given inequality $\sum_i \Delta v_i(\eta, \lambda_i) \leq \bar{z} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\|$. If λ^δ is bounded for every δ then $\{\lambda^\delta\}$ has a convergent subsequence with limit λ^0 . But this λ^0 violates the conditions of [Theorem 4](#), so assume that $\{\lambda^\delta\}$ is unbounded. A deviation plan σ_i^r is called *relatively undetectable* if $\sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| = 0$ whenever $\eta \in \Delta(A)$ satisfies $\sum_i \Delta v_i(\eta, \lambda_i) \leq \bar{z} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\|$. Call \mathcal{D}_i^r the set of relatively undetectable plans. A deviation plan σ_i^s is called *relatively detectable* if

$$\max_{(p, \sigma_i, \sigma_i^r)} \{p : p\sigma_i^r + (1-p)\sigma_i = \sigma_i^s, \sigma_i \in \mathcal{D}_i^r, \sigma_i^r \in \mathcal{D}_i^r, p \in [0, 1]\} = 0.$$

Let \mathcal{D}_i^s be the set of relatively detectable plans. By the same argument as for [Lemma B.7](#), \mathcal{D}_i^s is a compact set, therefore, by the same argument as for [Lemma B.8](#),

$$\gamma_i^s := \min_{\sigma_i^s \in \mathcal{D}_i^s} \max_{\eta \in \Delta(A)} \left\{ \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \sigma_i^s)\| : \sum_{i \in I} \Delta v_i(\eta, \lambda_i) \leq \bar{z} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i)\| \right\} > 0.$$

Without loss, $\lambda_i^\delta = \lambda_i^{r,\delta} + \lambda_i^{s,\delta}$, where $\lambda_i^{r,\delta}$ is relatively undetectable and $\lambda_i^{s,\delta}$ is relatively detectable. By assumption, $\lambda_i^{r,\delta}$ is μ -unprofitable, so $\sum_{(b_i, \rho_i)} \lambda_i^{s,\delta}(a_i, b_i, \rho_i)$ is bounded below by $\beta > 0$, say. (Otherwise, $\sum_i \Delta v_i(\mu, \lambda_i^\delta) < \varepsilon$ for small $\delta > 0$.) But this implies that

$$\max_{\eta \in \Delta(A)} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i^\delta)\| = \max_{\eta \in \Delta(A)} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \lambda_i^{s,\delta})\| \geq \beta \gamma_i^s > 0.$$

But this contradicts our initial assumption, which establishes the result. \square

References

- ABREU, D., P. MILGROM, AND D. G. PEARCE (1990): “Information and Timing in Repeated Partnerships,” *Econometrica*, 59, 1713–33. [22](#)
- ALCHIAN, A. AND H. DEMSETZ (1972): “Production, Information Costs, and Economic Organization,” *American Economic Review*, 62, 777–795. [2](#), [3](#), [7](#), [19](#), [20](#), [21](#), [24](#)
- AOYAGI, M. (2005): “Collusion Through Mediated Communication in Repeated Games with Imperfect Private Monitoring,” *Economic Theory*, 25, 455–475. [22](#), [23](#)
- CHE, Y.-K. AND J. KIM (2006): “Robustly Collusion-Proof Implementation,” *Econometrica*, 74, 1063–1107. [17](#)
- CHESLOW, D. (2005): “Students, Bartenders Cited in Liquor Sting,” *The Daily Northwestern*, <http://www.dailynorthwestern.com/>. [3](#)

- CREMER, J. AND R. MCLEAN (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1257. 17
- D’ASPREMONT, C. AND L.-A. GÉRARD-VARET (1998): “Linear Inequality Methods to Enforce Partnerships under Uncertainty: An Overview,” *Games and Economic Behavior*, 25, 311–336. 2, 11, 22
- FORGES, F. (1986): “An Approach to Communication Equilibria,” *Econometrica*, 54, 1375–1385. 4, 9
- FUCHS, W. (2007): “Contracting with Repeated Moral Hazard and Private Evaluations,” *American Economic Review*, 97, 1432–1448. 2, 7, 20, 21
- FUDENBERG, D., D. LEVINE, AND E. MASKIN (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039. 11, 22
- HOLMSTRÖM, B. (1982): “Moral Hazard in Teams,” *Bell Journal of Economics*, 13, 324–340. 2, 5, 6, 20
- KANDORI, M. (2003): “Randomization, Communication, and Efficiency in Repeated Games with Imperfect Public Monitoring,” *Econometrica*, 71, 345–353. 22
- KNIGHT, F. H. (1921): *Risk, Uncertainty, and Profit*, Boston, MA: Schaffner & Marx, Houghton Mifflin Company. 24
- LEGROS, P. AND H. MATSUSHIMA (1991): “Efficiency in Partnerships,” *Journal of Economic Theory*, 55, 296–322. 2, 22
- LEGROS, P. AND S. MATTHEWS (1993): “Efficient and Nearly Efficient Partnerships,” *Review of Economic Studies*, 60, 599–611. 2, 22
- LEHRER, E. (1992): “On the Equilibrium Payoffs Set of Two Player Repeated Games with Imperfect Monitoring,” *International Journal of Game Theory*, 20, 211–226. 23
- LELAND, H. AND H. PYLE (1977): “Informational Asymmetries, Financial Structure, and Financial Intermediation,” *Journal of Finance*, 32, 371–387. 24
- LEVIN, J. (2003): “Relational Incentive Contracts,” *American Economic Review*, 93, 835–847. 2, 7, 20, 21
- MACLEOD, B. (2003): “Optimal Contracting with Subjective Evaluation,” *American Economic Review*, 93, 216–240. 2, 20
- MACLEOD, B. AND J. MALCOMSON (1998): “Motivation and Markets,” *American Economic Review*, 88, 388–411. 21

- MYERSON, R. (1986): “Multistage Games with Communication,” *Econometrica*, 54, 323–358. 4, 9
- (1997): “Dual Reduction and Elementary Games,” *Games and Economic Behavior*, 21, 183–202. 32, 33
- NAU, R. F. AND K. F. MCCARDLE (1990): “Coherent Behavior in Noncooperative Games,” *Journal of Economic Theory*, 50, 424–444. 32, 33
- OBARA, I. (2008): “Folk Theorem with Communication,” *Journal of Economic Theory*, forthcoming, mimeo. 22
- PRENDERGAST, C. (1999): “The Provision of Incentives in Firms,” *Journal of Economic Literature*, 37, 7–63. 2, 20, 21
- RADNER, R., R. MYERSON, AND E. MASKIN (1986): “An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria,” *Review of Economic Studies*, 53, 59–69. 21
- RAHMAN, D. (2005a): “Optimum Contracts with Public and Private Monitoring,” Mimeo. 1
- (2005b): “Team Formation and Organization,” Ph.D. dissertation, UCLA. 1
- RAHMAN, D. AND I. OBARA (2008): “Mediated Partnerships,” Mimeo. 21
- ROCKAFELLAR, R. T. (1970): *Convex Analysis*, Princeton, New Jersey: Princeton University Press. 25
- STRAUSZ, R. (1997): “Delegation of Monitoring in a Principal-Agent Relationship,” *Review of Economic Studies*, 64, 337–357. 2, 20
- TADELIS, S. (2002): “The Market for Reputations as an Incentive Mechanism,” *Journal of Political Economy*, 110, 854–882. 21
- TOMALA, T. (2009): “Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring,” *Games and Economic Behavior*, to appear. 22, 23
- TSA (2004): “Guidance on Screening Partnership Program,” *Transportation Security Administration*, see also www.tsa.gov/what_we_do/screening/covert_testing.shtm. 3