

But Who will Monitor the Monitor?*

David Rahman[†]

University of Minnesota

September 6, 2010

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? In this paper I propose a contract that makes the monitor responsible for the monitoring technology, and thereby successfully provides incentives even when the monitor's observations are not only private, but costly, too. I also characterize exactly when such a contract can provide monitors with the right incentives to perform. In doing so, I emphasize virtual enforcement and suggest its implications for the theory of repeated games.

JEL Classification: D21, D23, D82.

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

*Alchian and Demsetz (1972, p. 782).

[†]Financial support from the Spanish Ministry of Education's Research Grant No. SEJ 2004-07861 while at Universidad Carlos III de Madrid as well as the National Science Foundation's Grant No. SES 0922253 is gratefully acknowledged. An early version of this paper was circulated under the title "Optimum Contracts with Public and Private Monitoring," which was based on Chapter 3 of my Ph.D. dissertation at UCLA. I owe many thanks to Antonio Cabrales, V. V. Chari, Harold Demsetz, Andrew Dust (for excellent research assistance), Willie Fuchs, Larry Jones, Narayana Kocherlakota, David Levine, Roger Myerson, Ichiro Obara (whose collaboration on a related paper spilled over into this one), Joe Ostroy, Bill Zame and numerous seminar audiences for insightful comments that helped me tremendously.

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 airport inspectors that sit behind an X-ray machine, watching suitcases pass them by. Their "output" is paying attention—only they can know if they are scrutinizing the baggage in front of them or just daydreaming. Of course, this problem is closely related to that of providing incentives for security guards and regulatory agencies, as well as maintaining police integrity. Without the right incentives, these agents might be tempted to shirk on their responsibilities or succumb to possibly unchecked corruption. Naturally, this problem appears in many other economic realms, such as management and supervision of workers, especially in service industries.

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, I 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 then will Bob know what to tell Ann in order to receive his wage. Ann can now rest assured that Bob will be honest and obedient.

Contrived though it may seem, this kind of contract is ubiquitous. The Transportation Security Administration uses “covert testing” to evaluate airport inspectors (TSA, 2004, p. 5). Such testing ranges from superimposing images of bombs on computer screens to smuggling weapons. In 2005, the Government Accountability Office created a “Forensic Audits and Special Investigations Unit” (GAO, 2007, p. 11). This unit has undertaken several “red team” operations to expose vulnerabilities in government agencies, including the Federal Emergency Management Agency, the Nuclear Regulatory Commission, the Department of Defense, the Department of Transportation, Medicare, and the U.S. Department of Labor, to name a few. These operations have ranged from smuggling nuclear materials through the U.S. border to test the effectiveness of border patrols, to making unfair labor practice claims to test the Wage and Hour Division’s reporting standards, and obtaining Medicare billing numbers without proper documentation to test due diligence in Medicare’s protocol for granting billing rights (GAO, 2008a,b,c,d, 2009a,b).

Similar arrangements for “policing the police” are also well-documented. Internal Affairs Departments regularly use “integrity tests” to discourage police corruption.³ Officers are

³Sherman (1978, pp. 163–4) on integrity tests by police departments: *Both Oakland and New York constructed artificial situations giving police officers the opportunity to commit corrupt acts. The tests were designed to yield the evidence needed to arrest and convict an officer who failed the “test.” [...] Some were random tests of conformity of procedures, such as the infamous “wallet drop”: wallets containing marked money were dropped by internal policing officers near randomly selected patrol officers to see if they would turn the wallet in to the police property clerk with the full amount of money. Other integrity tests had more*

infiltrated into police corruption rings to act as informants. In both cases, the mere possibility of monitoring can deter corruption.⁴ Even corrupt officers have been used to “test” and report on other officers, often relying on leniency to provide incentives.^{5,6} Two important examples of the investigation and disciplining methods above according to the criminology literature are the Knapp Commission (Knapp et al., 1972) and the Mollen Commission (Mollen, 1994) investigations. See Marx (1992) for more interesting examples.

Comparable contracts have also been used by managers. For instance, retailers routinely hire “mystery shoppers” (Ewoldt, 2004) to secretly evaluate employees and provide feedback to managers. (Airlines call them “ghost riders,” see Dallos, 1987.) The consulting branch of IBM offers “ethical hacking” services by “tiger teams” (Palmer, 2001) that try to hack into clients’ IT network, to expose vulnerabilities. Amazon’s Mechanical Turk (Pontin, 2007) decentralizes a wide variety of tasks to humans, such as verifying image quality. To provide workers with incentives, images whose quality is already known are occasionally included.

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 question of how to remunerate monitors,⁷ it has generated much academic debate. Previously, a monitor’s observations were assumed to be *specific targets*. *Money left in an illegally parked car was often used to test the integrity of certain police tow-truck drivers against whom allegations of stealing from towed cars had been made. Fake gambling operations were created to see if police officers tried to establish paid protection arrangements.*

⁴Sherman (1978, pp. 156–7) on informants: *Under careful monitoring, honest police officers in New York were even assigned by internal investigators to join corruption conspiracies [...]. Quite apart from the value of the information these regular informants provided, the very fact that their existence was known to other police officers may have yielded a deterrent effect. Though the informants were few in number, no one was quite certain who could be trusted to keep silence.* See also Prenzler (2009, pp. 137–8) for several examples.

⁵Sherman (1978, p. 162): on incentives for informants: *[...] rewards were used to encourage informants to inform. These ranged from immunity from arrest for the informants’ own offenses to simple obligation for future considerations.* See also Skolnick (1994, pp. 112–138).

⁶This is close to the DOJ’s Corporate Leniency Program, which, to discourage collusion, awards immunity to the first firm in a cartel to come forward with evidence of illegal activity. (Harrington, 2008; Miller, 2009.)

⁷Juvenal asked a similar question (see Hurwicz, 2008) 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.

either publicly verifiable (Footnote 1) or costless (Footnote 2). This paper adds to the debate by constructing a theoretical model that accommodates *costly private monitoring*. Existing solutions from the literature fail to provide the right incentives in this richer environment, but in this paper I show how to make the monitor responsible for monitoring with a version of Bob’s contract, and characterize exactly when this contract is enforceable.

I begin my analysis (Section 1) by studying a firm with two agents: a worker (Friday) and a monitor (Robinson). I show how the principal can align incentives by allocating different private information to each agent and making Robinson’s reward depend on both his report and Friday’s information. Formally, I consider contracts that form a *communication equilibrium* (Forges, 1986; Myerson, 1986) as follows. Usually 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 he worked, whereas if Friday was asked to shirk then Robinson will be paid only if he reports back that he 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, but if he monitors he’ll observe Friday’s behavior and deduce the owner’s effort recommendation, which secures his payment. Therefore, every agent is honest and obedient in equilibrium.

For my next main result, I study a general environment and reconcile the following infinite regress inherent to monitoring.⁸ Suppose that providing incentives for workers requires costly private monitoring to detect their deviations. What about the monitor’s deviations? Theorem 4 characterizes when the monitor’s deviations can be discouraged. First I show (Theorem 1) that discouraging deviations one by one is enough to discourage them simultaneously. I then argue that detectable deviations can be discouraged with a contract similar to Bob’s. On the other hand, an undetectable deviation still detects workers’ deviations by virtue of being undetectable—since by definition it is (statistically) indistinguishable from monitoring no matter what anyone does—so it continues to fulfill a monitoring role. Now

⁸A related regress on monitors for monitors is studied by Basu, Bhattacharya and Mishra (1992).

let the monitor play this deviation. Of course, this argument also applies to the monitor’s deviations from this deviation, and so forth. This infinite regress (of the monitor playing a deviation of a deviation of ...) is reconciled by showing that under reasonable assumptions (e.g., if agents have finitely many choices) not every behavior by the monitor has a profitable, undetectable deviation. Thus, to induce workers’ effort, their deviations must be detectable with occasional monitoring, but deviations from monitoring need not be detectable.

There are many ways to interpret this result. One is in terms of [Becker’s \(1968\)](#) model of crime and punishment. Even if having neither crime nor enforcers is an impossible ideal, it may be approached with very little crime and very few enforcers (but large penalties to criminals) if and only if every crime is detectable in the weak sense alluded to above. Another interpretation is as the weakest detectability requirement for the Folk Theorem and related results in the theory of repeated games. I discuss these connections in [Sections 3 and 4](#).

1 Robinson and Friday

Example 1. 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.⁹ Finally, after Robinson observes the realized signal, he makes a verifiable report to the principal.

⁹Alternatively, we could assume that if Robinson rests he observes “no news.” The current assumption helps to compare with the literature that relies on publicly verifiable monitoring, such as [Holmström \(1982\)](#).

If monitoring were costless—following the subjective evaluation literature (Footnote 2)—the principal could enforce the action profile (monitor,work) by paying Robinson a wage independent of his report. Robinson would be willing to monitor and report truthfully, and Friday could therefore be rewarded contingent on his effort via Robinson’s 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. Unfortunately, this is impossible, since if Robinson rests then Friday’s expected payment cannot depend on his own effort, so he will shirk. On the other hand, if Robinson’s observations are publicly verifiable then not only can the principal enforce (monitor,work), but also *virtually enforce* (rest,work)—i.e., enforce an outcome arbitrarily close—using Holmström’s group penalties: if news is good everyone gets paid and if news is bad nobody gets paid. Thus, 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 .

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, although Robinson would happily tell the truth with a wage independent of his report, he would never monitor, so again Friday would shirk. This begs the question: How can we motivate Friday to work when Robinson’s signal is both 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 correlating Robinson’s payment with Friday’s secret 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 asked to monitor with probability σ , (ii) Friday is independently asked to work with probability μ , and (iii) the principal pays Robinson and Friday, respectively, contingent on his recommendations and Robinson’s report as follows.

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

Robinson and Friday’s recommendation- and report-contingent payments

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. Robinson is not told Friday's recommendation—this he must discover by monitoring. Clearly, Friday is willing to obey the principal's recommendations if Robinson is honest and obedient. To see that Robinson will abide by the principal's requests, suppose that he was asked to monitor. If he monitors, clearly it is optimal for him to also be honest, with expected payoff $\mu(1/\mu) + (1-\mu)[1/(1-\mu)] = 2$. If instead he rests, 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 tends to the profile (rest,work) with a contract that makes the behavior incentive compatible along the way. In other words, (rest,work) is *virtually enforceable*. This requires arbitrarily large payments, and in reality feasible payments may be bounded. (For more on this, see [Section 3.2](#).) Nevertheless, virtual enforcement is a useful benchmark describing what outcomes are approachable with sufficiently large payments. Interpreting payments as continuation values in a dynamic game, virtual enforcement also describes what is attainable as agents become sufficiently patient.

The profile (rest,work) is virtually enforced by having Friday shirk with positive probability. With public monitoring, (rest,work) was virtually enforced by incurring the cost of monitoring Friday (Robinson's effort) with small probability. With private monitoring, an additional cost is incurred, also with small probability: the cost of monitoring Robinson. This cost is precisely the foregone productivity from Friday shirking. Such a loss may be avoided by asking Friday to take a costless action, like changing the color of his socks.

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 through his ability to reproduce Friday's recommendation. As such, Robinson 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 simply told Robinson his recommendation. However, this is cheap talk—it still is an equilibrium that they don't share this information. (For more on this, see [Section 3.3](#).)

This example shows that a monitor need not be the principal in order for his incentives to be aligned, contrary to [Alchian and Demsetz's](#) claim, who would argue for making Robinson the principal. Indeed, if Robinson was the principal then he would never verify Friday's effort, as Friday's payment would come from his own pocket *after* effort had 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."¹⁰ 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.

A final comment: if recommendations are not verifiable, $(\text{rest}, \text{work})$ is still virtually enforceable without a third party by asking Friday if he worked. [Section 4.2](#) has the details.

2 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 $\mu \in \Delta(A)$.¹¹ 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 signal profiles. A *monitoring technology* is a measure-valued map $\text{Pr} : A \rightarrow \Delta(S)$, where $\text{Pr}(s|a)$ stands for the probability that $s \in S$ was observed given that $a \in A$ was played.

¹⁰Several 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 some incentive for the principal to under-report effort that cannot be overcome dynamically. As a result, some inefficiency remains.

¹¹If 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 .

Incentives are provided with linear transfers. An *incentive scheme* is any given function $\zeta : I \times A \times S \rightarrow \mathbb{R}$ that assigns individual payments contingent on *recommended* actions and *reported* signals, each of which is assumed verifiable. (I relax verifiability in Section 4.2.)

Time elapses as follows. First, the principal commits to a *contract* (μ, ζ) , draws a profile of suggestions according to μ , and delivers them to the agents confidentially and verifiably.¹² Agents now simultaneously take unverifiable and unobservable actions. Next, agents observe their private, unverifiable signals and submit a report before a public signal realizes (the order of public and private signals is not essential for the results, 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 $\rho_i : S_i \rightarrow \bar{S}_i$, where $\rho_i(s_i)$ is the reported signal when agent i observes s_i . Let R_i be the set of i 's reporting strategies. The *truthful reporting strategy* is the identity map $\tau_i : S_i \rightarrow \bar{S}_i$ 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 is reported if everyone else is honest and plays a_{-i} equals

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

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

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

In other words, (μ, ζ) is incentive compatible if μ is a *communication equilibrium* (Myerson, 1986; Forges, 1986) of the game induced by ζ .

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 $\mu^m \rightarrow \mu$.

¹²Since the principal can commit to any contract, I equate the principal and the mediator without loss of generality. I discuss the consequences of relaxing this assumption in Section 4.

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 everyone is honest and obedient except for agent i , who plays σ_i instead, defined for each signal profile s by

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

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 $s \in S$.¹³ 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 recommendation profile in B such that the report probabilities induced by the strategy differ from that induced by honesty and obedience, assuming that everyone else is honest and obedient.

I begin with an intuitive characterization of enforceability. Given 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 . (Each strategy set is clearly convex.) The principal pays the surrogate the expected deviation gains from i playing σ_i instead of honest and obediently,

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

Clearly, the value of this game is at least zero for the surrogate, since he could always have his agents play honest and obediently. In fact, by construction, μ is enforceable if and only if this value equals zero, since the value is zero if and only if there is an incentive scheme that discourages every strategy by the surrogate. 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.

¹³We abuse notation by identifying Dirac measure $[a] \in \Delta(A)$ with the action profile $a \in A$.

Now pick any strategy σ_i . If it is $\text{supp } \mu$ -detectable¹⁴ then there exists an action profile $a \in \text{supp } \mu$ 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 utility gain 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 strategy 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. Call these strategies *deviations*. 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*.

Theorem 2. *Fix any correlated strategy μ . Every $\text{supp } \mu$ -disobedience is $\text{supp } \mu$ -detectable if and only if for any profile of utility functions, μ is enforceable.*

¹⁴Let $\text{supp } \mu = \{a \in A : \mu(a) > 0\}$ be the set of action profiles with positive probability under μ .

[Theorem 2](#) gives an intuitive characterization of robust enforceability¹⁵ of an outcome: every disobedience must be detectable with behavior in the support of the outcome. It is important to note that 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, Levine and Maskin \(1994\)](#).¹⁶ To illustrate, consider an example where every disobedience is detectable but IFR fails.¹⁷

Example 2. Two publicly verifiable signals and two agents, Ann and Bob. Ann has two choices, $\{U, D\}$, and Bob has three, $\{L, M, R\}$. The monitoring technology is given 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 then Bob playing R is statistically indistinguishable from $\frac{1}{2}[L] + \frac{1}{2}[M]$. Similarly, if Ann plays D then Bob can deviate from R to play $\frac{1}{3}[L] + \frac{2}{3}[M]$ without changing signal probabilities. Hence, IFR fails. In [Section 3.1](#), I show that it is impossible to get Bob to ever play R with an incentive scheme that only depend on signals if Bob strictly prefers playing L and M . However, every disobedience is detectable: for any deviation by Bob there is an action by Ann that detects it. By correlating Bob’s payment with Ann’s (recommended) action, the principal can keep Bob from knowing how 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 μ appears in [Theorem 2](#). This leads to the next result.

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.

¹⁵I.e., for all utility functions. For more on robustness, see [Section 3.3](#).

¹⁶IFR implies that the same correlated strategy detects every deviation. See [Footnote 21](#) for a definition.

¹⁷It even fails *local* IFR of [d’Aspremont and Gérard-Varet \(1998\)](#), which requires that one correlated strategy—possibly different for different agents—detect all of that agent’s deviations.

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

I will now characterize virtual enforcement of a correlated strategy, rather than every one, so 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 condition fails, as the next example shows.

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 utility profile, since $\{(U, L)\}$ -undetectable $\{(U, L)\}$ -disobediences exist, 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 $\{(U, L)\}$ -disobediences. The key condition here is that every $\{(U, L)\}$ -disobedient deviation is detectable.

Theorem 4. *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 4 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 perhaps occasional behavior—call it “monitoring.” Crucially, there is no requirement on disobediences to behavior outside of μ , so deviations from monitoring need not be detectable.

To make intuitive sense of this result, 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 $\sigma_j(a_j)$? After all, the condition of Theorem 4 purposely says nothing about detection outside B .

If such $\sigma_j(a_j)$ is detectable then it can easily be discouraged. 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 intuition completes my answer to the question “But who will monitor the monitor?” The principal monitors the monitor’s detectable deviations by occasionally asking his workers to secretly shirk, and nobody needs to monitor the monitor’s undetectable deviations.¹⁸ 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. The monitor can be made responsible with contractual terms that follow Robinson’s incentive scheme in [Example 1](#). These terms provide monitors with incentives for reporting accuracy.

Since [Theorem 4](#) only depends on the support of μ , [Theorem 3](#) and the next result follow.

Corollary 2. *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.*

To end this section, I characterize virtual enforcement for a fixed utility profile. Although exact enforcement has a simple characterization ([Theorem 1](#)), a corresponding result for virtual enforcement is not as simple. To see why, notice that, on the one hand, virtually enforcing some μ does not require that every supp μ -disobedience be detectable: unprofitable supp μ -disobediences may be undetectable, for instance. On the other hand, it is not enough that every profitable supp μ -disobedience be detectable, as the next example shows.

Example 4. Consider the following variation on Robinson and Friday ([Example 1](#)).

	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		

¹⁸This argument relies on mediated contracts, as they imply that different actions may be used to detect different deviations. Without them, enforcement requires that the same behavior detect every disobedience, so for virtual enforcement, just what is undetectable changes along the way. See [Section 3.1](#) for details.

Assume that signals are *publicly verifiable* and Robinson’s utility is constant. Clearly, the profile (rest,work) is not enforceable, since Friday shirking is (rest,work)-profitable and {(rest,work)}-undetectable. Moreover, (rest,work) is *not virtually enforceable* either. Indeed, for Friday to ever work Robinson must monitor with positive probability. But then nothing can discourage Friday from playing solitaire, since it is statistically indistinguishable from working and weakly dominant. On the other hand, every (rest,work)-profitable disobedience is detectable: every (rest,work)-profitable strategy involves shirking with positive probability, and shirking is detectable.

Detecting (rest,work)-profitable deviations is not enough here because solitaire weakly dominates work and is indistinguishable from it. Indeed, if solitaire strictly dominated work then there would exist a (rest,work)-profitable, undetectable strategy, rendering (rest,work) virtually unenforceable. On the other hand, 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.

So what *is* required beyond detecting profitable deviations? Below, I will argue that profitable deviations must be *uniformly and credibly* detectable. To illustrate, note that if solitaire is removed from [Example 4](#) then (rest,work) is virtually enforceable, not just because every (rest,work)-profitable deviation is detectable (this is true with or without solitaire), but also because the utility gains from every (rest,work)-profitable deviation can be uniformly outweighed by monetary losses. To describe this “uniform detection” formally, 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 *supp μ -undetectable* if and only if $\|\Delta \Pr(\mu, \sigma_i)\| = 0$.

Say that *every μ -profitable deviation is uniformly detectable* if $z \geq 0$ exists such that for every μ -profitable deviation σ_i there is a correlated strategy η (possibly different for different σ_i) with σ_i being *supp η -detectable* and $\Delta v_i(\eta, \sigma_i) < z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_i)\|$. Uniform detectability says that a bound $z \geq 0$ exists such that for every μ -profitable deviation σ_i , there exist (i) a correlated strategy η that detects σ_i , and (ii) an incentive scheme ζ satisfying

$-z \leq \zeta_i(a, s) \leq z$, that strictly discourage σ_i .^{19,20} Intuitively, every μ -profitable deviation can be strictly discouraged with a correlated strategy and an incentive scheme bounded by the same amount. To see how uniform detectability fails in [Example 4](#) but holds without solitaire, let α and β , respectively, be the probabilities that Friday shirks and plays solitaire after being asked to work. Clearly, (rest,work)-profitability requires $\alpha > 0$. To get uniform detectability we need z such that given $\alpha > 0$ and $\beta \geq 0$, a correlated strategy η exists with $(\alpha + \beta)\eta(\text{monitor,work}) + \alpha\eta(\text{rest,work}) < 2z\alpha\eta(\text{monitor,work})$. Therefore, $(\alpha + \beta)/\alpha < 2z$ is necessary for uniform detectability. However, no z satisfies this for all relevant (α, β) . Removing solitaire restores uniform detectability: now $\beta = 0$, so any $z > 1/2$ works.

Uniform detectability is still not enough for virtual enforcement, as [Example 5](#) shows.

Example 5. Add a row to the table in [Example 4](#), i.e., another action for Robinson, with utility payoffs

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

 and signal probabilities

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

.

In [Example 5](#) every (rest,work)-profitable deviation is uniformly detectable when Robinson plays his new action, but this action is not credible because it is strictly dominated by and indistinguishable from monitoring. Hence, (rest,work) is not virtually enforceable.

Definition 4. Say that *every μ -profitable deviation is uniformly and credibly detectable* if there exists $z \geq 0$ such that for every μ -profitable deviation σ_i , there exists a correlated strategy η satisfying (i) σ_i is $\text{supp } \eta$ -detectable, (ii) $\Delta v_i(\eta, \sigma_i) < z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_i)\|$, and (iii) $\Delta v_j(\eta, \sigma_j) \leq z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_j)\|$ for all other (j, σ_j) .

Intuitively, just as before, we may use different η to uniformly detect different σ_i , but these η must be credible in that incentives can be provided ([Footnote 19](#)) for everyone else to play η . Finally, uniform and credible detection yields the characterization we sought.

Theorem 5. *A correlated strategy μ is virtually enforceable if and only if every μ -profitable deviation is uniformly and credibly detectable.*

¹⁹To find this payment scheme, invoking the Bang-Bang Principle, let $\zeta_i(a, s) = \pm z$ depending on the sign of the statistical change created by σ_i , namely $\sum_{(b_i, \rho_i)} \sigma_i(b_i, \rho_i | a_i) \Pr(s | a_{-i}, b_i, \rho_i) - \Pr(s | a)$.

²⁰To see intuitively why I have a strict inequality, suppose that there is a μ -profitable deviation σ_i and a correlated strategy η for which $\Delta v_i(\eta, \sigma_i) = z \sum_a \eta(a) \|\Delta \Pr(a, \sigma_i)\|$. In this case, it would be impossible to discourage σ_i by playing η instead of μ with some probability and payments bounded within $\pm z$.

3 Discussion

In this section I discuss possible extensions and limitations of the model. I begin by describing the contractual margin of recommendation-contingent payments and noting how [Theorem 4](#) relies deeply on them. Next, I study enforceability subject to bounded payments as well as other contractual restrictions. Finally, I discuss robustness, collusion and genericity.

3.1 The Margin of Mediated Contracts

The results of [Sections 1](#) and [2](#) rely crucially on mediated rewards: incentive schemes that depend on the principal’s recommendations. [Examples 1](#) and [2](#) show that such schemes can yield a strict improvement for the principal relative to ones that just depend on reported signals. I now characterize intuitively this improvement. I suggest a way to interpret this contractual enrichment as relaxing the detectability requirements of enforcement. Intuitively, I will argue that recommendation-contingent schemes effectively allow for deviations to be detected “after the fact” in the sense that different actions may be used to detect different deviations. I will then show how this relates to [Theorem 4](#) in important ways.

Given μ , call σ_i *detectable at μ* if $\Pr(s|\mu) \neq \Pr(s|\mu, \sigma_i)$ for some $s \in S$.²¹

Theorem 6. *Fix any correlated strategy μ . Every μ -profitable deviation is detectable at μ if and only if μ is enforceable with an incentive scheme that is independent of recommendations.*

[Theorem 6](#) captures the potential of mediated contracts relative to incentive schemes that do not depend on recommendations. By [Theorem 2](#), this potential is characterized by the difference between supp μ -detectability and detectability at μ . To interpret the difference, fix a correlated strategy μ and 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 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 utility profile. In other words, with recommendation-contingent rewards it is as if the

²¹This definition is very close to (and implied by) individual full rank (IFR). Formally, IFR (at μ) means that every deviation σ_i is detectable at μ , although σ_i need not be a non-negative vector.

principal chooses a correlated strategy *after* the surrogate chooses a disobedience, in order to detect it. To illustrate, recall [Example 1](#). If Robinson is asked to monitor but instead chooses to rest and report g then the principal can react by asking Friday to shirk, which would have led to b had 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, without recommendation-contingent rewards it is as if the principal chooses μ first in the hide and seek game and subsequently the surrogate looks for a disobedience undetectable at μ . Thus, Robinson monitoring is not enforceable without recommendation-contingent payments. (See also [Example 2](#) for similar logic.)

For instance, consider enforcing a pure strategy profile a . By [Theorem 1](#), this requires that every a -profitable disobedience be $\text{supp } [a]$ -detectable. Of course, $\text{supp } [a]$ -detectability coincides with detectability at a . Since agents receive only one recommendation under a , there is no use for incentive schemes that depend on recommendations. However, when enforcing a correlated strategy with non-singleton support, the two contract spaces differ, and as such so do the appropriate notions of detectability that characterize enforcement. In general, even though both with and without mediated contracts an allocation μ is enforceable if and only if every profitable deviation is detectable (either $\text{supp } \mu$ - or at μ), with mediated contracts different actions (in the support of μ) may be used to detect different deviations, whereas without them μ itself must simultaneously detect every profitable deviation.

This shows how [Theorem 4](#) relies deeply on recommendation-contingent rewards. To conclude that a monitor’s undetectable deviations are irrelevant, we argued that any such deviation is just as good because it does not change signal probabilities after any action profile. Hence, the set of undetectable deviations is independent of the correlated strategy being enforced, and remains unchanged at every approximation stage of virtual enforcement. Without recommendation-contingent rewards the relevant notion of detectability changes to “at μ .” Now the set of undetectable deviations depends on the correlated strategy being enforced, which changes at every approximation stage. Thus, a deviation from monitoring may be undetectable at one stage of the approximation but not at another. As a result, without mediated contracts a comparable version of [Theorem 4](#) must fail.

3.2 Other Contractual Restrictions

In this paper I make various simplifying assumptions, like ignoring limited liability, budget balance and individual rationality. I now discuss how these assumptions affect the results.

For instance, the contracts of [Example 1](#) and the characterizations of virtual enforcement in [Theorems 4](#) and [5](#), use arbitrarily large payments. Although many important contributions to contract theory rely on them, such as [Becker \(1968\)](#) and [Holmström \(1982\)](#), it is also of interest to understand the consequences of removing this arguably unrealistic assumption.

Imposing one-sided limited liability on agents does not change the paper’s results simply because a constant can be added to any payment scheme without disrupting its incentive properties. Therefore, an allocation is (virtually) enforceable if and only if it is so subject to agents’ one-sided limited liability. Of course, limited liability will generally affect *optimal* contracts. However, in this paper I am not looking for optimal contracts but rather feasible ones. On the other hand, two-sided limited liability will restrict the set of enforceable outcomes if monetary losses cannot be made sufficiently large to discourage otherwise profitable deviations. Nevertheless, the spirit of [Theorems 1](#) and [6](#) remains: deviations only need to be discouraged one by one and different actions can be used to detect different deviations, although now the amount of detection must outweigh utility gains from a deviation.

Theorem 7. *A correlated strategy μ is enforceable with payments bounded above and below by $\pm z$ if and only if $\Delta v_i(\mu, \sigma_i) \leq z \sum_a \mu(a) \|\Delta \Pr(a, \sigma_i)\|$ for each deviation σ_i .*

Interestingly, restricting the incentive scheme in [Theorem 7](#) to not depend on recommendations yields $\Delta v_i(\mu, \sigma_i) \leq z \|\Delta \Pr(\mu, \sigma_i)\|$ as a characterizing condition. This shows how μ must simultaneously detect every deviation by an amount that outweighs its utility gain (again, [Footnote 19](#) applies). In any case, virtual enforcement is clearly subject to two-sided limited liability. Mathematically, the set of virtually enforceable correlated strategies is the closure of the set of enforceable ones. With two-sided limited liability, the set of enforceable outcomes is already closed, so virtual and exact enforcement coincide. Of course, as liability limits become relaxed (i.e., $z \rightarrow \infty$), the set of enforceable outcomes approaches that of virtually enforceable ones in the sense of [Theorems 4](#) and [5](#).

Thus, consider an allocation that is virtually enforceable but not enforceable. With agents' liability limited by z , this allocation is not enforceable, but there is a "closest" one (in some meaningful sense) that is. As $z \rightarrow \infty$, this closest allocation approaches the given one. This motivates interpreting virtual enforcement as a benchmark that describes which outcomes are approachable before knowing agents' liability limits. Similarly, the study of virtual enforcement can help to understand the gain from relaxing these limits. An example of this comes from [Becker \(1968\)](#), who suggested the trade-off between the probability of monitoring and the size of contingent payments. [Theorems 4 and 5](#) may be interpreted as yielding necessary (and, of course, sufficient) conditions for [Becker's](#) contracts to provide the right incentives, and [Theorem 7](#) as characterizing the trade-off between the probability of playing a monitoring action and the size of contingent payments.

As for limits on the principal's liability, it is possible that, ex post, the principal cannot pay an aggregate amount larger than some upper bound. This constraint, together with one-sided limited liability on the agents, also affects the set of feasible outcomes. (On its own, it clearly does not affect the feasible set.) Nevertheless, similar results to those in [Section 2](#) may be obtained. To illustrate, this is what [Theorem 1](#) would become:

Theorem 8. *A correlated strategy μ is enforceable with an incentive scheme ζ that satisfies (i) limited liability for the agents, i.e., $\zeta_i(a, s) \leq 0$, and (ii) limited liability for the principal, i.e., $\sum_i \zeta_i(a, s) \geq -z$ for some given $z \geq 0$, if and only if*

$$\sum_{i \in I} \Delta v_i(\mu, \sigma_i) + z \sum_{(a, s)} \mu(a) \min_{i \in I} \{\Delta \Pr(s|a, \sigma_i)^-\} \leq 0 \quad (\dagger)$$

for each profile $(\sigma_1, \dots, \sigma_n)$ of strategies, where $\Delta \Pr(s|a, \sigma_i)^- = \min\{\Delta \Pr(s|a, \sigma_i), 0\}$.

The proof of [Theorem 8](#) is omitted, as it is just a technical complication of that for [Theorem 1](#). Let us briefly interpret the result. Given a profile of strategies, the first term of (\dagger) stands for the sum of expected unilateral deviation gains across agents arising from this profile, and the second term stands for the z -weighted expected maximal probability decrease across agents (think of this decrease as good news). The principal's liability constraint relates each agent's incentive scheme, so it is no longer without loss to provide incentives agent by agent. However, as in [Theorem 1](#), it suffices to discourage strategy profiles one by one. For

any such profile, whoever decreases the probability of a signal profile the most after each recommendation profile is rewarded the maximum amount z —everyone else is paid nothing. An outcome is thus enforceable with limited liability on the principal and agents if and only if, after this arrangement, every strategy profile is unprofitable in the sense of (†). Imposing limited liability on both principal and agents implies that the set of feasible incentive schemes lies in a compact set. Therefore, virtual and exact enforcement coincide here, too. Finally, even if different individuals have different liability limits, the standard result (see Legros and Matsushima, 1991, for instance) still applies: only total liability matters for enforceability.

Regarding budget balance, on the one hand, ex ante budget constraints do not affect any results because a constant may be added to any incentive scheme so as to satisfy them. On the other hand, ex post budget balance—that the sum of payments across agents is fixed at every state of the world—does have a substantive impact on the set of enforceable outcomes. This issue is addressed by Rahman and Obara (2010), who characterize enforceable allocations subject to budget balance. They show that, in addition to detecting disobediences, identifying obedient agents is necessary (and sufficient) for budget-balanced enforcement. However, that paper fails to find necessary and sufficient conditions for virtual enforcement—with or without budget balance. See their paper for details.

As for participation constraints, by themselves they also do not bind for enforceability, again by adding a constant to any incentive scheme that satisfies all such constraints. In conjunction with other constraints such as the ones discussed above, they will generally bind. Rahman and Obara (2010) consider the effect of participation constraints together with budget balance, and Rahman (2010) extends these results to an arbitrary setting.

As a final comment, although adding some or all of the above constraints may add realism to the model, I view them as layers on top of the basic constraints of incentive compatibility. Broadly, the main themes of this paper survive these complications, such as versions of the Minimax Lemma and the margin of recommendation-contingent rewards. Although, as discussed, in some cases virtual enforcement collapses to exact, in other important cases it still does not. Even when it does collapse, virtual enforcement serves as a useful benchmark for the value of relaxing the constraints that preclude it.

3.3 Robustness

I end this section by discussing how the results of the paper are robust. I will focus on robustness with respect to (i) preferences,²² (ii) agents’ beliefs, (iii) the monitoring technology, and (iv) renegotiation. Broadly, I argue below that mediated contracts can enforce outcomes robustly and the conditions of some the results in the paper under which enforcement is robust are themselves robust.

Of course, a result like [Theorem 1](#) cannot be robust to the set of profitable, detectable deviations because it is an “if and only if” statement. However, the point of the theorem is to characterize enforceability, so its objective is to find weak conditions. As a result, the theorem sacrifices robustness in favor of fine-tuning. On the other hand, some other results may be viewed as robust with respect to several notions of robustness, as I argue next.

Since everything is finite in the model, agents’ utility gains from a deviation are bounded. Without loss, normalize this bound to one. [Theorem 2](#) yields the following consequence.

Corollary 3. *Fix any correlated strategy μ . Every supp μ -disobedience is supp μ -detectable if and only if there is an incentive scheme ζ such that*

$$\mathbf{1}_{\{a_i \neq b_i\}} \leq \sum_{(a_{-i}, s)} \mu(a) \zeta_i(a, s) [\Pr(s|a_{-i}, b_i, \rho_i) - \Pr(s|a)] \quad \forall (i, a_i, b_i, \rho_i).$$

Therefore, the same incentive scheme ζ implements μ for any profile of utility functions $v : I \times A \rightarrow \mathbb{R}$ such that $\sup_{(i, a, b_i)} \{v_i(b_i, a_{-i}) - v_i(a)\} \leq 1$.

[Corollary 3](#) shows just how [Theorem 2](#) is robust to the utility profile: *there exists a single incentive scheme that simultaneously discourages every disobedience regardless of preferences* (as long as deviation gains are bounded and this bound is known). Clearly, the argument applies also to [Theorem 4](#). Therefore, the results of the paper that quantify with respect to the utility profile may be viewed as robust in the above sense. [Corollary 3](#) says more: there is a payment scheme that makes honesty and obedience a “strict equilibrium,” i.e., all disobedience-related incentive constraints hold with strict inequality.²³

²²The case of risk aversion is mentioned by [Rahman \(2005\)](#) and discussed recently by [Strausz \(2010\)](#).

²³To see this, scale the scheme in [Corollary 3](#) enough to strictly outweigh the gains from any disobedience. To also strictly discourage dishonesties, just replace “disobedience” with “deviation” in the result.

The possibility of strict equilibrium yields robustness with respect to just about everything in the model. Thus, as long as the conditions of [Corollary 3](#) hold, there exists an incentive scheme that enforces a given allocation incentives even if an agent’s interim beliefs about others’ actions are slightly perturbed, and even if the monitoring technology is slightly perturbed.²⁴ This contrasts the surplus-extracting contracts of [Cremer and McLean \(1985, 1988\)](#), for instance. Since they extract all surplus, participation constraints must hold with equality, which makes their contracts vulnerable to small changes in fundamentals, i.e., preferences and beliefs. I discuss them further in [Section 4.1](#).

There is another measure of robustness to the monitoring technology. As long as there are enough action-signal pairs for every agent’s opponents, I now argue that every disobedience is detectable generically on the set of monitoring technologies, i.e., except for those in a set of Lebesgue measure zero. Intuitively, incentives may be provided to an agent in three ways: (a) using only others’ signals to detect deviations, (b) using only own reports and others’ recommendations, and (c) using both one’s reports and others’ signals together. [Theorem 9](#) below identifies conditions such that for every agent, at least one such way is generic.

Theorem 9. *Every disobedience is detectable generically if for every agent i ,*

- (a) $|A_i| - 1 \leq |A_{-i}| (|S_{-i}| - 1)$ when $|S_i| = 1$,²⁵
- (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 $|S_{-i}| > 1$.

Intuitively, genericity holds even if $|S| = 2$, as long as agents have enough actions. Hence, a group of agents may overcome their 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.²⁶ Thus, if the monitoring technology were chosen uniformly at random then almost surely every disobedience would be detectable, so by [Theorem 3](#) every correlated strategy would be virtually enforceable.

²⁴By “slightly perturbed” I mean “within some neighborhood.”

²⁵In comparison, genericity for IFR with public monitoring requires much more: $|A_i| \leq |S| - 1$ for every i .

²⁶I thank Roger Myerson for urging me to emphasize this point.

I end this section by briefly discussing collusion, multiple equilibria, and renegotiation.

A notable weakness of the contracts in this paper is their overt failure to be collusion-proof. Thus, in our leading example (Section 1) Robinson and Friday could communicate “extra-contractually” and break down the incentives that the secret recommendations tried to provide.²⁷ On the other hand, collusion is a general contractual problem.²⁸ For instance, the surplus-extracting scheme of Cremer and McLean (1988) is not collusion-proof for similar reasons. In this paper I have not tried to overcome collusion because I view it as an additional layer, although it is possible to find conditions under which collusion amongst agents may be overcome with mediated contracts.²⁹ For instance, one may restrict attention to constant-surplus contracts, as in Che and Kim (2006).

Since the contracts of this paper rely on communication, they are subject to equilibrium multiplicity: there always exist “babbling” equilibria where agents ignore their recommendations if everyone else does. This precludes “full implementation.” (See also Kar, Ray and Serrano, 2010.) Equilibrium multiplicity brings with it the usual baggage of equilibrium selection, but this baggage may not be all bad: it is also a babbling equilibrium to ignore extra-contractual messages, which restores some resilience of mediated contracts to collusion.

Finally, renegotiation may also be included as an additional layer to the problem outlined in this paper, to obtain conditions for mediated contracts that are robust in this sense. Some forms of renegotiation have been likened in the literature to efficiency and budget balance (see, e.g., Neeman and Pavlov, 2010, and references therein). See Rahman and Obara (2010) characterize allocations that are enforceable with budget balanced payments.

²⁷The following incentive scheme deters such communication between Robinson and Friday (Friday weakly 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	$0, 0$	$1/(1 - \mu), 0$	$0, 1/(1 - \sigma)$	$1/2(1 - \mu), 1/2(1 - \sigma)$

However, it is impossible to prevent Robinson from buying Friday’s signal. On the other hand, in keeping with Alchian and Demsetz’s original motivation, if there were another worker (say, Saturday) then it would be possible to provide the right collusion-proof incentives (see Section 4.1 for more).

²⁸The likelihood of collusion depends on agents’ ability to communicate, so it need not always be a problem.

²⁹See the working paper version of this paper, Rahman (2009).

4 Literature

In this section I compare the results of this paper with the relevant literature. I begin with a discussion of the partnership problem, which motivated [Alchian and Demsetz](#) as well as others to study the role of monitoring, together with the literature on subjective evaluation and mechanism design. Finally, I discuss several papers in the literature on contracts and repeated games with comparable detectability criteria and continuation values.

4.1 The Partnership Problem

[Alchian and Demsetz](#)'s partnership problem 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?

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.*³⁰ 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.³¹ Hence, making the monitor residual claimant, or principal, fails to provide the right incentives.

³⁰*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.*

³¹A 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.

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 shows 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,³² and as such [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 importance 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 provide 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. However, 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? In this case, *no matter how small the cost*, rewarding a monitor independently of his report induces him to avoid exerting effort towards reporting accurately.

In this paper I accommodate costly private monitoring and find a contract that gives both workers and monitors the right incentives to perform. It also addresses the partnership problem. Although [Example 1](#) had just one worker (Friday), it is easy to add one more, call him Saturday. Following [Alchian and Demsetz \(1972\)](#), suppose that, if an 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

³²Intuitively, 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.

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 \(2010\)](#) for related work that characterizes enforceability with budget balance.³³

Virtual enforcement was explored in some detail by [Legros and Matthews \(1993\)](#). Using mixed strategies, they found sufficient conditions for virtually enforcing an efficient outcome with budget-balanced payments, so-called "nearly efficient partnerships." Their sufficient conditions are stricter than those in [Rahman and Obara \(2010\)](#), mainly because they relied on output-contingent contracts. The latter paper fleshes out the differences in detail. [Miller \(1997\)](#) enriches the model of [Legros and Matthews](#) by adding costless private monitoring. As a result, his contracts are subject to similar criticisms as those of subjective evaluation.

The partnership problem has also been addressed from a dynamic perspective (e.g., [Radner, Myerson and Maskin, 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, I abstract from such dynamic instruments but acknowledge that they are implicit in the paper. This interpretation is discussed further in [Section 4.2](#).

There is also 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, market incentives may be incorporated via participation constraints.

³³[Rahman and Obara \(2010\)](#) finds sufficient but not necessary conditions for virtual enforcement.

Let me now discuss some mechanism design theory. The literature on surplus extraction, developed by [Cremer and McLean \(1985, 1988\)](#) and [McAfee and Reny \(1992\)](#), relies on exogenous correlation in agents' private information to discipline reporting. This correlation enables the principal to extract agents' surplus. An extreme but instructive example is two agents with perfectly correlated types. If their reports differ, then the principal knows that someone lied, so he punishes everyone. Thus, if one agent tells the truth then the other agent will, too, so truth-telling is an equilibrium. There are similarities with mediated contracts, but also important differences. Firstly, types are correlated endogenously in my model: the principal allocates private information to provide incentives. Secondly, I don't always need every agent's report to provide incentives. For instance, with Robinson and Friday the principal told Friday his type rather than soliciting it from him, whereas [Cremer and McLean](#) solicit information from every agent. Thirdly, I focus on enforceability rather than surplus extraction, which is clearly weaker. In this sense, the conditions of the paper are less restrictive than those for surplus extraction, so my contracts are more general as a result. Finally, as I argued in [Section 3.3](#), my contracts are robust in a way that theirs need not be.

The work of [Mezzetti \(2004, 2007\)](#) is also related. He observes that when agents' values are interdependent, their realized utilities are correlated conditional on the public outcome (here the outcome is the action profile). Hence, it is possible to discipline agents further by conditioning their payments on reported utility profiles. In a sense, this paper generalizes his results by viewing utility realizations as monitoring signals.

4.2 Detection and Enforcement

Several papers have emphasized the duality between detection and enforcement. Some of the first ones to point this out are [Abreu, Milgrom and Pearce \(1990\)](#) and [Fudenberg, Levine and Maskin \(1994\)](#), in the context of repeated games. Together with the 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, and none of them considers recommendation-contingent rewards. Therefore, they fail to enforce many outcomes that are enforceable in this paper, even in a more general context.

Some recent papers have studied rich contract spaces in specific settings, like [Kandori \(2003\)](#) and its private monitoring version by [Obara \(2008\)](#), [Aoyagi \(2005\)](#) and [Tomala \(2009\)](#). [Aoyagi](#) uses dynamic mediated strategies that rely on “ ε -perfect” monitoring, and fail if monitoring is costly or one-sided. [Tomala](#) studies recursive communication equilibria and independently considers recommendation-contingent continuation values to prove a folk theorem. He derives a version of the Minimax Lemma, but does not consider virtual enforcement. [Kandori](#) has agents play mixed strategies and report the realization of their mixtures. He studies contracts contingent on those reports and signal realizations. Mediated contracts can outperform [Kandori’s](#) even with public monitoring, as the next example shows.

Example 6. One agent, three actions (L , M and R), two publicly verifiable signals (g and b), and 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 mediated contracts but not with [Kandori’s](#). 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 supposedly plays σ and is then asked what he played before getting paid. He gains two ‘utils’ by playing M instead and reporting L (R) if the realized signal is g (b), with the same expected monetary payoff.

Theorem 10. *If agents secretly report their actions before they observe any signals then [Kandori’s](#) contracts generate the same set of enforceable mixed strategy profiles as mediated contracts. If not, mediated contracts generally dominate [Kandori’s](#).*

[Theorem 10](#) clarifies the difference between [Kandori’s](#) and mediated contracts, and suggests an improvement to [Kandori’s](#) contracts: have agents mix and report their *intended* action (as in pool tables everywhere) before taking it. Now the principal’s recommendations need not be verifiable, but then reported intentions must be. Thus, in [Example 1](#) the principal can monitor [Robinson](#) by having [Friday](#) mix and report what he played.³⁴ Indeed, let σ

³⁴Interestingly, [Kandori’s](#) contracts suffice to enforce [Robinson](#) monitoring with probability one, but not with probability strictly between zero and one. Intended actions must be reported for this.

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 (rest,work), Robinson must either report what he plans to play before he plays it or what he played before he observes the signal. Here, the principal’s role is just being a budget breaker, i.e., there is no need for a mediator. If Robinson and Friday can commit to destroy or sell value, they can write this contract by themselves.

On the other hand, these “unmediated” contracts are not without fault. Firstly, without a mediator, enforcement is restricted to mixed strategy profiles, which may be undesirable. Thus, in the classic Chicken game (see, e.g., [Aumann, 1974](#)), maximizing welfare involves correlated equilibria that are not even public randomizations over Nash equilibria. Secondly, as agents mix in these contracts they must be made indifferent, whereas with mediated contracts they may be given strict incentives. Thus, robustness—as described in [Section 3.3](#) and in [Bhaskar’s \(2000\) critique](#) (see also [Bhaskar, Mailath and Morris, 2008](#)), for instance—fails with these other contracts, but holds with mediated ones. Finally, virtually enforcing even a pure strategy profile may require correlated strategies, as the next example shows.

Example 7. There are three agents—Rowena picks a row, Colin a column, and Matt a matrix. Rowena and Colin are indifferent over everything. Here is Matt’s utility function.

	L	R		L	R		L	R
U	1	2	U	0	0	U	-1	2
D	2	-1	D	0	0	D	2	1
	A			B			C	

There are two publicly verifiable signals. The monitoring technology is below.

	L	R		L	R		L	R
U	$\frac{1}{2}, \frac{1}{2}$	1, 0	U	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{2}$	U	$\frac{1}{2}, \frac{1}{2}$	0, 1
D	1, 0	1, 0	D	$\frac{1}{2}, \frac{1}{2}$	$\frac{1}{2}, \frac{1}{2}$	D	0, 1	0, 1
	A			B			C	

Clearly, the profile (U, L, B) is not enforceable, since playing A instead of B is a (U, L, B) -profitable, (U, L, B) -undetectable deviation. To virtually enforce (U, L, B) , Rowena and

Colin cannot play just (U, L) . But then playing $\frac{1}{2}[A] + \frac{1}{2}[C]$ instead of B is Matt's only undetectable deviation. Call this deviation σ . If only Rowena mixes and plays U with probability $0 < p < 1$ then σ is profitable: Matt's profit equals $2(1-p) > 0$. Similarly, if only Colin mixes between L and R then σ is profitable. If Rowena and Colin mix independently, with $p = \Pr(U)$ and $q = \Pr(L)$, Matt still profits from σ : he gets $2p(1-q) + 2(1-p)q > 0$. On the other hand, the correlated strategy $r[(U, L, B)] + (1-r)[(D, R, B)]$ for $0 < r < 1$ renders σ unprofitable, which is still Matt's only undetectable deviation. Therefore, (U, L, B) is virtually enforceable (by letting $r \rightarrow 1$), although only with correlated behavior.

Let me finish by discussing some other papers on repeated games, starting with [Compte \(1998\)](#) and [Kandori and Matsushima \(1998\)](#). These papers study games where players each period first take actions, then observe signals, and finally communicate. Apart from the repeated game context, their models are close to the present paper, except that here actions may be correlated, and reactions contingent on previous recommendations. The main result of these papers is a Folk Theorem with public communication. However, their conditions for enforcement are much stronger because they do not consider mediated continuation values. [Theorem 4](#) gives necessary and sufficient conditions for virtual enforceability, which may be interpreted as the limit set of outcomes when $\delta \rightarrow 1$. Also related is the work by [Kandori and Obara \(2006\)](#) and [Ely, Hörner and Olszewski \(2005\)](#) on belief-free equilibria, and [Kandori \(2009\)](#) on weakly belief-free equilibria. These papers consider continuation values that depend on others' actions and observations. However, they all rule out communication: continuation values cannot depend on one's own intended/recommended action *and* reported signal. Thus, none of these other papers can virtually enforce (rest,work) in [Example 1](#).³⁵

Finally, the work of [Lehrer \(1992\)](#) is especially noteworthy. He characterizes the equilibrium payoff set of a two-player repeated game with imperfect monitoring and time-average utilities (heuristically, discount factors equal one) 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. This is established by having players undertake detecting behavior as the game proceeds with probability diminishing so quickly that it does not affect time-average utility.

³⁵A similar comment applies to [Phelan and Skrzypacz \(2008\)](#) and [Kandori and Obara \(2010\)](#).

Without time-average utility, [Lehrer](#)'s argument fails. To characterize equilibrium payoffs as the discount factor δ tends to one rather than at the limit of $\delta = 1$, virtual enforcement is the appropriate notion, not sustainability. According to [Lehrer](#)'s result, the profile (rest,work) is sustainable in [Example 4](#) because every (rest,work)-profitable disobedience is detectable. However, as was already argued, it is not virtually enforceable. Understanding how [Theorem 5](#) helps to describe the limit of equilibrium payoffs in a repeated game as δ tends to one—especially the discontinuity of the equilibrium correspondence with respect to δ at $\delta = 1$ (see, e.g. [Radner, Myerson and Maskin, 1986](#))—is the object of future research.

5 Conclusion

In this paper I offer the following answer to [Alchian and Demsetz](#)'s question of who will monitor the monitor: The principal monitors the monitor's detectable deviations by having his workers occasionally shirk, and nobody needs to monitor the monitor's undetectable deviations ([Theorem 4](#)). How to monitor the monitor? With "trick questions," as in [Robinson's](#) contract ([Example 1](#)). This contract makes the monitor responsible for monitoring.

[Alchian and Demsetz](#) argued that the monitor must be made residual claimant for his incentives to be aligned. In a sense, they "elevated" the role of monitoring in organizations. On the other hand, I have argued for "demoting" their monitor 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 classic question is the purpose of future research.

[Knight \(1921, Part III, Ch. IX, par. 10\)](#) 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 a screening role for residual claims. Again, according to [Knight \(1921, Part III, Ch. IX, par. 11\)](#): *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.

References

- Abreu, Dilip, Paul Milgrom, and David G. Pearce.** 1990. "Information and Timing in Repeated Partnerships." *Econometrica*, 59(6): 1713–33. 28
- Alchian, A., and H. Demsetz.** 1972. "Production, Information Costs, and Economic Organization." *American Economic Review*, 62(5): 777–795. 0, 3, 8, 24, 25, 26, 32
- Aoyagi, Masaki.** 2005. "Collusion Through Mediated Communication in Repeated Games with Imperfect Private Monitoring." *Economic Theory*, 25: 455–475. 29
- Aumann, Robert.** 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics*, 1: 67–96. 30
- Basu, K., S. Bhattacharya, and A. Mishra.** 1992. "Notes on Bribery and the Control of Corruption." *Journal of Public Economics*, 48(3): 349–359. 4
- Becker, G.S.** 1968. "Crime and punishment: An economic approach." *Journal of Political economy*, 76(2). 5, 19, 20
- Bhaskar, V.** 2000. "The robustness of repeated game equilibria to incomplete payoff information." *University of Essex*. 30
- Bhaskar, V., G.J. Mailath, and S. Morris.** 2008. "Purification in the infinitely-repeated prisoners' dilemma." *Review of Economic Dynamics*, 11(3): 515–528. 30
- Che, Yeon-Koo, and Jinwoo Kim.** 2006. "Robustly Collusion-Proof Implementation." *Econometrica*, 74(4): 1063–1107. 24

- Compte, Olivier.** 1998. “Communication in Repeated Games with Imperfect Private Monitoring.” *Econometrica*, 66(3): 597–626. 31
- Cremer, Jacques, and Richard McLean.** 1985. “Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent.” *Econometrica*, 53(2): 345–361. 23, 28
- Cremer, Jacques, and Richard McLean.** 1988. “Full extraction of the surplus in Bayesian and dominant strategy auctions.” *Econometrica*, 56(6): 1247–1257. 23, 24, 28
- Dallos, Robert E.** 1987. “‘Ghost Riders:’ Airlines Spy on Selves in Service War.” *Los Angeles Times*, July 21. 3
- d’Aspremont, Claude, and Louis-André Gérard-Varet.** 1998. “Linear Inequality Methods to Enforce Partnerships under Uncertainty: An Overview.” *Games and Economic Behavior*, 25: 311–336. 1, 12, 28
- Ely, Jeffrey C., Johannes Hörner, and Wojciech Olszewski.** 2005. “Belief-Free Equilibria in Repeated Games.” *Econometrica*, 73(2): 377–415. 31
- Ewoldt, J.** 2004. “Dollars and Sense: Undercover Shoppers.” *Star Tribune*, October 27. 3
- Forges, F.** 1986. “An approach to communication equilibria.” *Econometrica*, 54(6): 1375–1385. 4, 9
- Fuchs, William.** 2007. “Contracting with Repeated Moral Hazard and Private Evaluations.” *American Economic Review*, 97(4): 1432–1448. 1, 8, 26, 27
- Fudenberg, Drew, David Levine, and Eric Maskin.** 1994. “The Folk Theorem with Imperfect Public Information.” *Econometrica*, 62(5): 997–1039. 12, 28
- GAO.** 2007. “GAO Strategic Plan 2007-2012.” Government Accountability Office, Washington, D.C. GAO-07-1SP. 2
- GAO.** 2008a. “Border Security: Summary of Covert Tests and Security.” Government Accountability Office, Washington, D.C. GAO-08-757. 2

- GAO.** 2008*b*. “Investigative Operations: Use of Covert Testing to Identify Security Vulnerabilities and Fraud, Waste, and Abuse.” Government Accountability Office, Washington, D.C. GAO-08-286T. 2
- GAO.** 2008*c*. “Medicare: Covert Testing Exposes Weaknesses in the Durable Medical Equipment Supplier Screening Process.” Government Accountability Office, Washington, D.C. GAO-08-955. 2
- GAO.** 2008*d*. “Undercover Tests Reveal Significant Vulnerabilities in DOT’s Drug Testing Program.” Government Accountability Office, Washington, D.C. GAO-08-225T. 2
- GAO.** 2009*a*. “Covert Testing Shows Continuing Vulnerabilities of Domestic Sales for Illegal Export.” Government Accountability Office, Washington, D.C. GAO-09-725T. 2
- GAO.** 2009*b*. “Wage and Hour Division Needs Improved Investigative Processes and Ability to Suspend Statute of Limitations to Better Protect Workers Against Wage Theft.” Government Accountability Office, Washington, D.C. GAO-09-629. 2
- Harrington, J.E.** 2008. “Optimal Corporate Leniency Programs.” *The Journal of Industrial Economics*, 56(2): 215–246. 3
- Holmström, B.** 1982. “Moral Hazard in Teams.” *Bell Journal of Economics*, 13: 324–340. 1, 5, 6, 7, 19, 26
- Hurwicz, L.** 2008. “But who will Guard the Guardians?” *American Economic Review*, 98(3): 577–585. 3
- Kandori, M.** 2009. “Chicken in Prison: Weakly Belief-Free Equilibria in Repeated Games with Private Monitoring.” University of Tokyo CIRJE Discussion Paper F-491. 31
- Kandori, Michihiro.** 2003. “Randomization, Communication, and Efficiency in Repeated Games with Imperfect Public Monitoring.” *Econometrica*, 71(1): 345–353. 29
- Kandori, Michihiro, and Hitoshi Matsushima.** 1998. “Private Observation, Communication, and Collusion.” *Econometrica*, 66(3): 627–652. 31

- Kandori, Michihiro, and Ichiro Obara.** 2006. "Efficiency in Repeated Games Revisited: The Role of Private Strategies." *Econometrica*, 74: 499–519. 31
- Kandori, Michihiro, and Ichiro Obara.** 2010. "Towards a belief-based theory of repeated games with private monitoring: An application of POMDP." Working Paper. 31
- Kar, A., I. Ray, and R. Serrano.** 2010. "A difficulty in implementing correlated equilibrium distributions." *Games and Economic Behavior*, 69(1): 189–193. 24
- Knapp, W., et al.** 1972. "Report of the Commission to Investigate Alleged Police Corruption." *New York: George Braziller.* 3
- Knight, Frank H.** 1921. *Risk, Uncertainty, and Profit.* Boston, MA: Schaffner & Marx, Houghton Mifflin Company. 32
- Legros, Patrick, and Hitoshi Matsushima.** 1991. "Efficiency in Partnerships." *Journal of Economic Theory*, 55(2): 296–322. 1, 21, 28
- Legros, Patrick, and Steven Matthews.** 1993. "Efficient and Nearly Efficient Partnerships." *Review of Economic Studies*, 60(3): 599–611. 1, 27, 28
- Lehrer, Ehud.** 1992. "On the Equilibrium Payoffs Set of Two Player Repeated Games with Imperfect Monitoring." *International Journal of Game Theory*, 20: 211–226. 31, 32
- Leland, H., and H. Pyle.** 1977. "Informational Asymmetries, Financial Structure, and Financial Intermediation." *Journal of Finance*, 32(2): 371–387. 33
- Levin, Jonathan.** 2003. "Relational Incentive Contracts." *American Economic Review*, 93(3): 835–847. 1, 8, 26, 27
- MacLeod, Bentley.** 2003. "Optimal Contracting with Subjective Evaluation." *American Economic Review*, 93(1): 216–240. 1, 26
- MacLeod, Bentley, and James Malcomson.** 1998. "Motivation and Markets." *American Economic Review*, 88(3): 388–411. 27

- Marx, G.T.** 1992. “When the guards guard themselves: Undercover tactics turned inward.” *Policing and Society*, 2(3): 151–172. 3
- McAfee, R.P., and P.J. Reny.** 1992. “Correlated Information and Mechanism Design.” *Econometrica*, 60(2): 395–421. 28
- Mezzetti, C.** 2004. “Mechanism design with interdependent valuations: Efficiency.” *Econometrica*, 72(5): 1617–1626. 28
- Mezzetti, C.** 2007. “Mechanism design with interdependent valuations: Surplus extraction.” *Economic Theory*, 31(3): 473–488. 28
- Miller, Nathan H.** 2009. “Strategic leniency and cartel enforcement.” *The American Economic Review*, 99(3): 750–768. 3
- Miller, Nolan H.** 1997. “Efficiency in Partnerships with Joint Monitoring.” *Journal of Economic Theory*, 77(2): 285–299. 27
- Mollen, M.** 1994. “Commission to investigate allegations of police corruption and the anti-corruption procedures of the police department: Commission report.” *City of New York, New York, NY*. 3
- Myerson, R.** 1986. “Multistage games with communication.” *Econometrica*, 54: 323–358. 4, 9
- Myerson, R.** 1997. “Dual Reduction and Elementary Games.” *Games and Economic Behavior*, 21(3): 183–202. 47
- Nau, R. F., and K. F. McCardle.** 1990. “Coherent Behavior in Noncooperative Games.” *Journal of Economic Theory*, 50: 424–444. 47
- Neeman, Z., and G. Pavlov.** 2010. “Renegotiation-proof mechanism design.” *UWO Department of Economics Working Papers*. 24
- Obara, Ichiro.** 2008. “The Full Surplus Extraction Theorem with Hidden Actions.” *The B.E. Journal of Theoretical Economics*, 8(1). 29

- Palmer, C.C.** 2001. “Ethical hacking.” *IBM Systems Journal*, 40(3): 769–780. 3
- Phelan, C., and A. Skrzypacz.** 2008. “Beliefs and private monitoring.” *mimeo*. 31
- Pontin, Jason.** 2007. “Artificial Intelligence, with Help from the Humans.” *New York Times*, March 25. 3
- Prendergast, Canice.** 1999. “The Provision of Incentives in Firms.” *Journal of Economic Literature*, 37(1): 7–63. 1, 26, 27
- Prenzler, T.** 2009. *Police corruption: preventing misconduct and maintaining integrity*. CRC. 3
- Radner, Roy, Roger Myerson, and Eric Maskin.** 1986. “An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria.” *Review of Economic Studies*, 53(1): 59–69. 27, 32
- Rahman, David.** 2005. “Team formation and organization.” Ph.D. dissertation, UCLA. 22
- Rahman, David.** 2009. “But Who Will Monitor the Monitor?” *mimeo*. 24
- Rahman, David.** 2010. “Detecting Profitable Deviations.” *mimeo*. 21
- Rahman, David, and Ichiro Obara.** 2010. “Mediated Partnerships.” *Econometrica*, 78(1): 285–308. 21, 24, 27
- Rockafellar, R. T.** 1970. *Convex Analysis*. Princeton University Press. 39
- Sherman, L.W.** 1978. *Scandal and reform*. University of California Press. 2, 3
- Skolnick, J.H.** 1994. *Justice without trial*. Macmillan. 3
- Strausz, R.** 1997. “Delegation of Monitoring in a Principal-Agent Relationship.” *Review of Economic Studies*, 64(3): 337–357. 1, 25
- Strausz, R.** 2010. “Mediated Contracts and Mechanism Design.” *mimeo*. 22

Tadelis, Steven. 2002. “The Market for Reputations as an Incentive Mechanism.” *Journal of Political Economy*, 110(4): 854–882. 27

Tomala, Tristan. 2009. “Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring.” *Games and Economic Behavior*, 67: 682–694. 29

TSA. 2004. “Guidance on Screening Partnership Program.” *Transportation Security Administration*, see also www.tsa.gov/what_we_do/screening/covert_testing.shtm. 2

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. □

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 (*). □

Theorem 3. It remains to prove necessity. If there is an undetectable disobedience σ_i then pick the utility function v_i so that it is strictly dominant to play σ_i rather than behaving honest and obediently. Now it is impossible for a_i to be played with positive probability, where a_i is any action where σ_i strictly dominates. □

Theorem 4. 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 5. See the end of [Appendix B](#). \square

Theorem 6. Fix any $\mu \in \Delta(A)$. By the Alternative Theorem, every μ -profitable deviation is detectable at μ if and only if a scheme $\zeta : I \times S \rightarrow \mathbb{R}$ exists such that for all (i, a_i, b_i, ρ_i) , $\sum_{a_{-i}} \mu(a) [v_i(a_{-i}, b_i) - v_i(a)] \leq \sum_{(a_{-i}, s)} \mu(a) \zeta_i(s) [\Pr(s | a_{-i}, b_i, \rho_i) - \Pr(s | a)]$, as required. \square

Theorem 7. Follows from [Lemma B.2 \(i\)](#). \square

Theorem 9. By [Lemma B.1](#), detectability of every disobedience implies *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 $\Pr(a_i, s_i) \Pr(a_{-i}, s_{-i}) = \Pr(s | a)$. 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 $|A_i| = |A_i| |S_i| \leq |A_{-i}| |S_{-i}| - (|A_{-i}| - 1) = |A_{-i}| (|S_{-i}| - 1) + 1$ for every i . Similarly, there are $|A_i| - 1$ redundant rows when $|S_{-i}| = 1$. Since the intersection of finitely many generic sets is generic, the result now follows. \square

Theorem 10. For the first claim, let $\mu = (\mu_1, \dots, \mu_n)$ be any mixed strategy profile. Clearly, if it is enforceable with Kandori's contracts then it is enforceable with mediated contracts, since the mediator can give agents the same incentive scheme with recommended actions instead of reported ones. Conversely, suppose that μ is enforceable with mediated incentive scheme ζ . I will now find an incentive scheme ξ such that (μ, ξ) is incentive compatible and $\sum_{a_{-i}} \mu_{-i}(a_{-i}) v_i(a) = \sum_{(a_{-i}, s)} \mu_{-i}(a_{-i}) \xi_i(a, s) \Pr(s | a)$ for all i and $a_i \in \text{supp } \mu_i$. Define $c_i(a_i) = \sum_{a_{-i}} \mu_{-i}(a_{-i}) v_i(a) - \sum_{(a_{-i}, s)} \mu_{-i}(a_{-i}) \zeta_i(a, s) \Pr(s | a)$ and $\xi_i(a, s) = \zeta_i(a, s) + c_i(a_i)$. This ξ enforces μ with Kandori's reports. This proves the first claim. The second claim now follows from the first together with [Example 6](#). \square

B Not for Publication: Ancillary Results

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. Otherwise, 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 $\Pr(s|a) = \sum_{(b_i, t_i)} \lambda_i(a_i, s_i, b_i, t_i) \Pr(s_{-i}, t_i|a_{-i}, b_i)$ for all (i, a, s) . Without loss, $\lambda_i(a_i, s_i, a_i, s_i) = 0$ for some (a_i, s_i) . To see this, note first that $\lambda_i(a_i, s_i, a_i, s_i) = 1$ for all (a_i, s_i) is impossible because $\sigma_i \geq 0$ is assumed disobedient. If $\lambda_i(a_i, s_i, a_i, s_i) \neq 1$, subtract $\lambda_i(a_i, s_i, a_i, s_i) \Pr(s|a)$ from both sides and divide by $1 - \lambda_i(a_i, s_i, a_i, s_i)$. Now $\Pr(a_i, s_i) \in \text{cone}\{\Pr(b_i, t_i) : (b_i, t_i) \neq (a_i, s_i)\}$ 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 $\max_{\sigma_i} \Delta v_i(\mu, \sigma_i) > 0$.³⁶

(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 .*

³⁶Intuitively, \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:

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

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.*

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

The correlated strategy μ is virtually 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 $z \geq 0$.

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

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

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

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

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*

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

Proof. The dual of this problem is given by

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

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

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

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 $\sum_{a_{-i}} \delta \mu(a) \Delta v_i(a, b_i) + \mathbf{1}_{\{b_i \neq a_i\}} \Delta v_i(a_i)^* \leq \sum_{a_{-i}} x_i(a) \cdot \Delta \text{Pr}(a, b_i, \rho_i)$ 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*

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

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

Call λ *extremely detectable* if for every (i, a_i) , $\lambda_i(a_i)$ cannot be written as a positive linear combination involving undetectable deviations. Let \mathcal{E} be the set of 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 $\forall(i, a_i), \sum_{a_{-i}} \Delta v_i(a_i, \sigma_i^e)^* - z \|\Delta \Pr(a, \sigma_i^e)\| \leq \|\Delta v\| - z \sum_{a_{-i}} \|\Delta \Pr(a, \sigma_i^e)\| \leq \|\Delta v\| - \frac{\|\Delta v\|}{\gamma} \gamma$.

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 $a \in A$,

$$\|\Delta \Pr(a, \lambda_j)\| = \|\Delta \Pr(a, \lambda_i^u, \lambda_j)\| \leq \|\Delta \Pr(a, \widehat{\lambda}_i^u, \lambda_j)\|,$$

where $\widehat{\lambda}_i^u(a_i, b_i, \tau_i) = \sum_{\rho_i} \lambda_i^u(a_i, b_i, \rho_i)$ and $\widehat{\lambda}_i^u(a_i, b_i, \rho_i) = 0$ 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 $\|\Delta \Pr(a, \lambda_i)\| \leq \|\Delta \Pr(a, \widehat{\lambda}_j^u, \lambda_i)\|$ 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 $\widehat{\lambda}_i^u$ -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 $\|\Delta \Pr(a, \lambda_i)\| \leq \|\Delta \Pr(a, \widehat{\lambda}_j^u, \lambda_i)\|$, 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. \square

Lemma B.11. *The conditions of Theorem 5 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 $\{\lambda^\delta\}$ is uniformly bounded then it has a convergent subsequence with limit λ^0 . But this λ^0 violates the conditions of Theorem 5, so $\{\lambda^\delta\}$ must be unbounded. Call a deviation σ_i^r *relatively undetectable* if given $\eta \in \Delta(A)$,

$\sum_i \Delta v_i(\eta, \sigma_i^r) \leq \bar{z} \sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \sigma_i^r)\|$ implies $\sum_{(i,a)} \eta(a) \|\Delta \Pr(a, \sigma_i^r)\| = 0$. Call \mathcal{D}_i^r the set of relatively undetectable plans. A deviation σ_i^s is called *relatively detectable* if

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

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

Proof of Theorem 5. 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 5 is fulfilled with $\eta = \mu$, 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 no m and m_1 exist with $\mu^m = p\mu^{m_1} + (1-p)\mu$ then μ^{m_2} exists 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 $\widehat{\mu}' \in B_\varepsilon(\widehat{\mu})$ such that $p\widehat{\mu} + (1-p)\mu = \widehat{\mu}'$ for some p such that $0 < p < 1$. Now, by the previous paragraph, the condition of [Theorem 5](#) 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, κ^z) solve $V_\mu(z)$ for every z . Given $\eta \in \Delta(A)$,

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

By the condition of [Theorem 5](#), \bar{z} exists with $\sum_{(i,a)} \Delta v_i(\eta^{\bar{z}}, \lambda_i^{\bar{z}}) < \bar{z} \sum_{(i,a)} \eta^{\bar{z}}(a) \|\Delta \Pr(a, \lambda_i^{\bar{z}})\|$ and $\sum_{(i,a)} \eta^{\bar{z}}(a) \|\Delta \Pr(a, \lambda_i^{\bar{z}})\| > 0$ for some $\eta^{\bar{z}}$, since $\lambda_i^{\bar{z}}$ is μ -profitable for some i . Hence, $C \leq 1 + (\bar{z} - z) \sum_{(i,a)} \eta^{\bar{z}}(a) \|\Delta \Pr(a, \lambda_i^{\bar{z}})\|$, i.e., $z - \bar{z} \leq (1 - C) / \sum_{(i,a)} \eta^{\bar{z}}(a) \|\Delta \Pr(a, \lambda_i^{\bar{z}})\|$. This inequality must hold for every z , therefore $\sum_{(i,a)} \eta^{\bar{z}}(a) \|\Delta \Pr(a, \lambda_i^{\bar{z}})\| \rightarrow 0$ as $z \rightarrow \infty$. But this contradicts [Lemma B.11](#), since $\sum_i \Delta v_i(\mu, \lambda_i^{\bar{z}}) \geq C$, completing the proof. \square