

Secret Contracts for Efficient Partnerships*

Ichiro Obara[†] and David Rahman[‡]

December 11, 2006

Abstract

By allocating different information to team members, secret contracts can provide better incentives to perform with an intuitive organizational design. For instance, they may help to monitor monitors, and attain approximately efficient partnerships by appointing a secret principal. More generally, secret contracts highlight a rich duality between enforceability and identifiability. It naturally yields necessary and sufficient conditions on a monitoring technology for any team using linear transfers to approximate efficiency (with and without budget balance). The duality is far-reaching: it is robust to complications in the basic model such as environments with infinitely many actions and signals. Thus, we obtain a subdifferential characterization of correlated equilibrium payoffs in discounted repeated games with imperfect private monitoring.

JEL Classification: D21, D23, D82.

Keywords: secret contracts, partnerships, duality, private monitoring.

*This paper subsumes the previously submitted “Optimum Contracts with Public and Private Monitoring,” by Rahman. We owe many thanks to Harold Demsetz, David Levine, Roger Myerson, Joe Ostroy, Phil Reny, Joel Sobel and Bill Zame for help with initial drafts of this paper. We also thank seminar audiences at UCLA, UC San Diego, University of Chicago, Universidad Carlos III de Madrid, Caltech, Universitat Pompeu Fabra, University of Minnesota, and University of Arizona.

[†]Department of Economics, University of California, Los Angeles.

[‡]Departamento de Economía, Universidad Carlos III de Madrid. Financial support from the Spanish Ministry of Education’s Research Grant No. SEJ 2004-07861 is gratefully acknowledged.

Contents

1	Introduction	1
1.1	Secrets and Monitors	2
1.2	Enforceability and Identifiability	3
2	Examples	5
2.1	Robinson and Friday	5
2.2	Secret Principal	7
3	Model	9
3.1	Linear Metering Problem	11
3.2	Unilateral Deviations	12
3.3	Obedient Players	15
4	Discussion	19
4.1	Exact versus Approximate Enforcement	19
4.2	Fixed Action Profiles and Utility Functions	20
4.3	Participation and Liability	21
4.4	Genericity	22
5	Extension	23
5.1	Equilibrium with Linear Programming	23
5.2	Continuous Metering Problem	26
5.3	Repeated Games with Private Monitoring	29
6	Conclusion	33
6.1	Abstract Mechanisms in Concrete Contracts	34
6.2	Secrets and Verifiable Recommendations	34
6.3	Usual Problems with Collusion	35
	References	35
A	Preliminaries	37
B	Proofs	38

1 Introduction

Ann owns a restaurant. She hires Bob to tally up the till every night and report back any mismatch between the till and that night's bills. Ann can motivate Bob to exert such effort and report truthfully any mismatch by secretly taking some money from the till herself with positive probability and offering him the following incentive scheme: 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 faces a secret contract: his report-contingent wage is unknown to him a priori (it depends on whether or not Ann secretly took some money). If Bob fails to exert effort, he won't know what to report in order to secure his wage. However, if he does his job he'll discover whether or not there is a mismatch and deduce from this Ann's behavior. Only after tallying the till will Bob know what to report in order to receive his wage, which turns out to be optimally truthful.

This paper studies contracts like Bob's and how they might help organizations to function productively. By allocating different information to team members, secret contracts often provide better incentives to perform with an intuitive organizational design. Thus, they give Bob incentives to acquire costly information and reveal it. In general, they provide a way of "monitoring the monitor" (Section 2.1), and can yield approximately efficient partnerships by appointing a "secret principal" (Section 2.2).

A rich duality between enforceability and identifiability—more specifically, between incentive compatible contracts and indistinguishable deviation plans—is exploited. It leads us to identify teams that can approximate efficiency (with and without budget-balanced transfers) by means of their "monitoring technology" (Section 3). This duality is far-reaching: it is amenable to complications in the basic model such as individual rationality and limited liability (Section 4). It also applies in environments with infinitely many actions and signals (Section 5). This last extension has useful implications, the most notable of which is perhaps a natural "subdifferential characterization" of dynamic equilibrium payoffs in discounted repeated games with imperfect private monitoring and private strategies (Section 5.3).

1.1 Secrets and Monitors

Monitoring is a central theme in this paper. 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.*¹ 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.²

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.* Where Alchian and Demsetz seem to overemphasize the role of monitoring in organizations, Holmström seems to underemphasize it. He provides incentives with “team punishments” that reward all agents when output is good and punish them all when it is bad. Assuming that output is publicly verifiable, he finds little role for monitoring,³ and perhaps as a result Holmström (1982, p. 339) concludes wondering: *... how should output be shared so as to provide all members of the organization (including monitors) with the best incentives to perform?*

Secret contracts motivate monitors: If the principal secretly recommends a worker to shirk or work, both with some probability (the worker can easily be motivated to willingly obey recommendations), and pays the monitor only if he reports back the recommendation, then—like Bob—the monitor will prefer to exert effort and report truthfully. To implement such contracts, the team requires (i) a disinterested mediator or machine that makes confidential, verifiable but non-binding recommendations to players, and (ii) transfers that depend on the mediator’s recommendation as well as the monitor’s report. As this requirement suggests, incentive compatibility of secret contracts is described here by Aumann’s (1974) *correlated equilibrium*.

¹*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.*

²A 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 verifiable. However, Strausz assumes that monitoring signals are “hard evidence,” so a monitor cannot misreport his information.

³Intuitively, if output were not publicly verifiable then his team punishments 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.

Monitoring adds value only insofar as it helps to provide incentives. Heuristically, if monitors never monitor then workers will not work, so costly monitoring may be worthwhile. Nevertheless, it is cost-efficient to do so as little as necessary. This leads naturally to *approximate efficiency* as the appropriate optimality criterion for a team with costly monitoring, especially when having access to *linear transfers*. For example, secret (mixed) monitoring of workers with small but positive probability together with large punishments if caught shirking saves costs while providing incentives.

This use of mixed strategies to approximate efficiency was developed by Legros and Matthews (1993) in Nash equilibrium with public, deterministic output. Not only can secret contracts exploit such mixing, too, but also (and in addition to monitoring the monitor) they can improve a team's contractual prospects even in the restricted setting of publicly verifiable output, as the secret principal demonstrates.

To see this, recall the partnership problem of Radner et al. (1986). It shows that no budget-balanced linear transfers contingent only on output can approximate efficiency in a team whose members can either work or shirk and whose joint output is (publicly verifiable and) either high or low with a probability that is increasing only in the number of workers. A secret principal approximates efficiency: With arbitrarily large probability, suppose everyone is recommended to work, and paid nothing regardless. With complementary probability, everybody is told to work except for one randomly picked team member, who is secretly told to shirk. This individual must pay everyone else if output is high and be paid by everyone else if output is low. Such a scheme is incentive compatible with large payments, budget-balanced, approximately efficient.

1.2 Enforceability and Identifiability

Assuming correlated equilibrium and approximate efficiency/enforceability renders linear our formal description of incentive compatible contracts. In other words, some given team behavior is approximately implementable with incentive compatible secret transfers if and only if a certain family of linear inequalities is satisfied. A duality theory of contracts therefore obtains as a result of this linearity, with basic implications for understanding incentives. We take advantage of this duality throughout the paper, which prevails over gradual complications in our basic model. Technically, our linear methods rely on Rahman (2005a) to extend those of Nau and McCardle (1990) and d'Aspremont and Gérard-Varet (1998) with substantially stronger results.

Duality yields two sides of the same coin, two opposite views of the same problem—in our case, a metering problem. As the title of this subsection—taken from Fudenberg et al. (1994, p. 1013)—suggests, enforceable contracts and unidentifiable deviation plans are mutually dual variables. As such, we obtain two natural descriptions of a team’s monitoring technology from each equivalent point of view. The primal side of the coin describes when contracts are enforceable and approximately implementable, whereas the dual side describes profiles of deviation plans that cannot be distinguished. Thus, the smaller the set of indistinguishable deviation plans, the larger the set of enforceable contracts—like a cone and its polar. In the limit, our main results (Theorems 3.6 and 3.10) identify necessary and sufficient conditions on a monitoring technology for any team outcome to be enforceable and approximately implementable via secret contracts (with and without budget balance).

Theorem 3.6 provides a minimal requirement on a team’s monitoring technology, called distinguishing unilateral deviations (DUD), that characterizes approximate enforceability with secret contracts of any team outcome. Intuitively, for every player there must be some opponents’ correlated strategy (not necessarily the same for everyone) that renders statistically identifiable disobeying the mediator. (Dishonesty may remain indistinguishable, though.) DUD turns out to be weak and generic.⁴

Restricting attention to budget-balanced secret contracts, Theorem 3.10 characterizes approximate enforceability of team behavior with a stronger condition, called identifying obedient players (IOP). Intuitively, IOP requires that—in addition to DUD—it is possible to statistically identify some player as obedient upon any deviation from some correlated strategy. IOP is weak⁵ and generic,⁶ too.

Our use of duality facilitates the study of other restrictions to the metering problem, like limited liability and individual rationality. Well-known results, such as that only total liability matters when providing a team with incentives or that reasonably low participation constraints don’t bind even with budget balance, are extended to this framework without complications. Exact implementation fits relatively nicely, too.

⁴DUD is weaker than comparable conditions in Compte (1998) and Obara (2006a). Restricted to public monitoring, DUD is weaker than local individual full rank, of d’Aspremont and Gérard-Varet (1998), which in turn is weaker than the condition in Legros and Matsushima (1991).

⁵IOP is weaker than comparable conditions such as those in Kandori and Matsushima (1998), Aoyagi (2005), and Tomala (2005). With public monitoring, it is still weaker than the compatibility of d’Aspremont and Gérard-Varet (1998) and even Kandori’s (2003) version of pairwise full rank.

⁶Like DUD, IOP is “as generic if not more” than other conditions in the literature (Section 4.4).

Teams with infinitely many actions and signals are also considered. Firstly, we characterize the set of correlated equilibria as well as their payoffs as the subdifferential of a value function in games with compact, Hausdorff action spaces and continuous utility functions. Building on the work of Hart and Schmeidler (1989), this result applies immediately to discounted repeated games with imperfect private monitoring, so as a corollary we characterize correlated equilibrium payoffs of any such game. We also prove a folk theorem with weak conditions (the weakest as far as we know) on a team’s monitoring technology. Finally, we extend the metering problem as well as our two main results to the infinite case. Everything (including DUD and IOP) generalizes mostly without complications, except for the need to reconcile infinitesimal deviations. To this end, we restrict the monitoring technology so that, intuitively, infinitesimal deviations remain unprofitable if non-infinitesimal ones are.

Further discussion of secret contracts, particularly their relation to the theory of mechanism design, their susceptibility to collusion, possible alternatives, and future research, is deferred to the conclusion (Section 6).

2 Examples

We begin our formal study of secret contracts with the development of two important, motivating examples mentioned in the introduction: monitoring the monitor, and the secret principal. The first example studies an environment involving contractual variations on a three-player game that attempts to typify the strategic interaction between a principal, an agent, and a monitor. The second example finds an intuitive way of attaining approximately efficient partnership with budget-balanced contracts.

2.1 Robinson and Friday

There are three players. The first is Robinson, who can either monitor or shirk. The second is Friday, who can either work or shirk. The third player is a so-called mediating principal, a disinterested party who makes recommendations and enforces contingent contractual payments. For simplicity, suppose the principal’s utility is constant regardless of the outcome of the game. Robinson (the row player) and Friday (the column player) interact according to the left bi-matrix below.

	work	shirk
monitor	2, -1	-1, 0
shirk	3, -1	0, 0

Utility Payoffs

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

Signal Probabilities

The action profile (shirk,work) is Pareto efficient, since Robinson finds monitoring costly and it does not intrinsically add value. However, this strategy profile is not incentive compatible by itself, since Friday always prefers to shirk rather than work.

The team’s monitoring technology is given by a set $S = \{g, b\}$ —so there are only two possible signals contingent upon which contracts may be written—together with the conditional probability system given by the right bi-matrix above. In words, if Robinson shirks then both signals are equiprobable, whereas if he monitors then the realized signal will accurately identify whether or not Friday worked. Contractual payments are assumed to be denominated in a private good (“money”) that enters players’ utility linearly with unit marginal utility.

Clearly, the efficient strategy profile (shirk,work) cannot be implemented.⁷ However, we can get arbitrarily close: When signals are publicly verifiable, the correlated strategy⁸ $\sigma[(\text{monitor}, \text{work})] + (1 - \sigma)[(\text{shirk}, \text{work})]$ can be implemented for any $\sigma \in (0, 1]$ with Holmström’s *team punishments*. For example, paying Robinson \$2 and Friday $\$1/\sigma$ if g and both players zero if b makes (shirk,work) approximately implementable.

If only Robinson observes the signal, and it is not verifiable, then for the principal to write signal-contingent contracts, he must first solicit the realizations from Robinson, who may in principle misreport them. Notice that now team punishments break down, since not only will Robinson always report g and shirk, but also Friday will shirk. Furthermore, if Robinson was rewarded independently of his report then although he would happily tell the truth, he would find no reason to monitor.

Another possibility is to have Friday mix between working and shirking. On its own, this strategy doesn’t change Robinson’s incentives to either lie or shirk. However, if the principal and Friday correlate their play without Robinson knowing when, it is possible to “cross-check” Robinson’s report, thereby “monitoring the monitor.”

⁷If Robinson shirks then no signal-contingent contract can compensate Friday more when working than shirking, since each signal carries the same probability regardless of Friday’s effort.

⁸As a matter of notation, let $[a]$ stand for Dirac measure (or the pure strategy profile a living in the space of correlated strategies) for any action profile a .

Specifically, the following correlated strategy is incentive compatible given $\mu \in (0, 1)$:
 (i) Robinson is told to monitor with probability σ (and shirk with probability $1 - \sigma$),
 (ii) Friday is independently told to work with probability μ (to shirk with $1 - \mu$), and
 (iii) the principal correlates his contractual strategy with players' recommendations:

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

The numbers on the left are Robinson's contingent payments, and those on the right are Friday's. Thus, 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.

It is easily checked that Friday has incentives to obey the mediator, and Robinson to monitor when asked as well as report truthfully. This contract also approximately implements the efficient outcome (shirk,work) by letting $\sigma \rightarrow 0$ and $\mu \rightarrow 1$.

Particularly distinguishing properties of this contract are firstly that Robinson does not directly observe the principal's recommendation to Friday, and secondly that Robinson has the incentive to monitor inasmuch as he is rewarded for reporting accuracy. Notice also that Robinson's report only confirms to the principal his recommendation to Friday. As such, the principal strips away Robinson's a priori informational advantage, which is why his surplus can be extracted. This contract allocates private information to approximate efficiency, showing that a team without asymmetric information may find it beneficial to create some as part of its organizational design.

A salient problem with this contract is that it is not robust to "collusion." If Friday could tell Robinson his recommendation then he would, so both players could save on effort. We do not deal with collusion formally in this paper, and delegate its discussion to the conclusion. On the other hand, there is no other way for Friday to work with positive probability; especially none without secret contracts.

2.2 Secret Principal

A team has n individuals. Each team member i can either work ($a_i = 1$) or shirk ($a_i = 0$). Let $c > 0$ be each individual's cost of effort. Effort is not observable. Output is publicly verifiable and can be either good (g) or bad (b). The probability of g equals $P(\sum_i a_i)$, where P is a strictly increasing function of the sum of efforts.

Radner et al. (1986) showed that in this environment there do not exist budget-balanced output-contingent linear transfers to induce everyone to work, not even approximately. One arrangement that is not approximately efficient but nevertheless induces most people to work is appointing Holmström’s principal. Call this player 1 and define transfers as follows. For $i = 2, \dots, n$ let $\zeta_i(g) = k$ and $\zeta_i(b) = 0$ be player i ’s output-contingent linear transfer, for some $k \geq 0$. Let player 1’s transfer equal

$$\zeta_1 = - \sum_{i=2}^n \zeta_i.$$

By construction, the budget is balanced. Everyone but player 1 will work if k is sufficiently large. However, player 1 has the incentive to shirk. This contract follows Holmström’s suggestion to the letter: Player 1 is a “fixed” principal who absorbs the incentive payments to all others by “breaking” everyone else’s budget constraint.

Allowing now for secret contracts, consider the following scheme. For any small $\varepsilon > 0$, a mediator asks every individual to work (call this event $\mathbf{1}$) with probability $1 - \varepsilon$. With probability ε , he picks some player i at random (with probability ε/n for all i) and asks him secretly to shirk, while telling all others to work (call this event $\mathbf{1}_{-i}$). For $i = 1, \dots, n$ let $\zeta_i(g|\mathbf{1}) = \zeta_i(b|\mathbf{1}) = 0$ be player i ’s contingent transfer if the mediator asked everyone to work. Otherwise, if player i was secretly told to shirk, for $j \neq i$ let $\zeta_j(g|\mathbf{1}_{-i}) = k$ and $\zeta_j(b|\mathbf{1}_{-i}) = 0$ be player j ’s transfer. For player i , let

$$\zeta_i = - \sum_{j \neq i} \zeta_j.$$

Clearly, this contract is budget-balanced. It is also incentive compatible. Indeed, if player i is recommended to work, incentive compatibility requires that

$$\frac{\varepsilon(n-1)}{n} P(n-1)k - c \geq \frac{\varepsilon(n-1)}{n} P(n-2)k,$$

which is satisfied if k is sufficiently large. If player i is asked to shirk, we require

$$-(n-1)P(n-1)k \geq -(n-1)P(n)k - c,$$

which always holds.

Therefore, this contract implements the efficient outcome with probability $1 - \varepsilon$ and a slightly inefficient outcome with probability ε . Since ε can be made arbitrarily small (by choosing an appropriate reward k), we obtain an approximately efficient partnership. The role of principal is not fixed here. It is randomly assigned with very small probability to make negligible the loss from having a principal.

3 Model

Let $I = \{1, \dots, n\}$ be a finite set of players, A_i a finite set of actions available to player $i \in I$, and $A = \prod_i A_i$ the space of action profiles. Actions are neither verifiable nor directly observable. A *correlated strategy* is any probability measure $\sigma \in \Delta(A)$. The profile of individual utilities over action profiles is captured by a map $v : I \times A \rightarrow \mathbb{R}$. We denote by $v_i(a)$ the utility to a player $i \in I$ from action profile $a \in A$.

The team's monitoring technology is described as follows. We begin with a family $\{S_j : j \in I \cup \{0\}\}$ such that S_i is a finite set of *private signals* observable only by individual member $i \neq 0$ and S_0 consists of *publicly verifiable* signals. Let

$$S := \prod_{j=0}^n S_j$$

be the product space of all observable signals. A *monitoring technology* consists of the space S together with a measure-valued map

$$\text{Pr} : A \rightarrow \Delta(S)$$

where $\text{Pr}(s|a)$ stands for the conditional probability that s will be observed by the players given that the team adopts action profile a . For every $s \in S$, suppose there exists $a \in A$ such that $\text{Pr}(s|a) > 0$.

We assume that the team has access to *linear transfers*. An *incentive scheme* is a map $\zeta : I \times A \times S \rightarrow \mathbb{R}$, interpreted as a contract that assigns money payments contingent on individuals, recommended actions, and *reported* signals. This formulation assumes that *recommendations are verifiable*.⁹

Instead of studying incentive schemes ζ directly, we will focus on *probability weighted* transfers, $\xi : I \times A \times S \rightarrow \mathbb{R}$. For any recommendation $a \in A$ with $\sigma(a) > 0$, we may think of ξ as solving

$$\xi_i(a, s) = \sigma(a)\zeta_i(a, s)$$

for some ζ . For any $a \in A$ with $\sigma(a) = 0$ and $\xi(a) \neq 0$, one may think of ξ as either arising from unbounded incentive schemes (i.e., $\zeta_i(a, s) = \pm\infty$) or as the limit of a sequence $\{\sigma_n \zeta_n\}$. This change of variables from ζ to ξ is explained in Section 4.1.

⁹This is without loss of generality: even if recommendations are not directly verifiable, we can let the players announce their recommendations as verifiable messages. As will be seen, almost every correlated strategy has an incentive scheme similar to that of Cremer and McLean (1988) which solicits recommendations at no cost.

The timing of team members' interaction runs as follows. Firstly, players agree upon some *contract* (σ, ξ) consisting of a correlated strategy σ together with probability weighted transfers ξ . Recommendations are drawn according to σ and made to players confidentially and verifiably by some machine. Players then choose an action and take it. Once actions have been adopted, the players observe their private signals and submit a verifiable report of their observations (given by an element of their personal signal space) before observing the public signal (not essential, just simplifying), after which recommendation- and report-contingent transfers are made according to ξ .

If every player obeys his recommendation and reports truthfully, the utility to Mr. i (before recommendations are actually made) from a given contract (σ, ξ) equals

$$\sum_{a \in A} v_i(a) \sigma(a) + \sum_{(a,s)} \xi_i(a, s) \Pr(s|a).$$

Of course, Mr. i may disobey his recommendation to play some action $b_i \in A_i$ and lie about his privately observed signal. A *reporting strategy* is a map $\rho_i : S_i \rightarrow S_i$, where $\rho_i(s_i)$ is the reported signal when Mr. i privately observes s_i . Let R_i be the set of all reporting strategies for player i . The *truthful reporting strategy* is the identity map $\tau_i : S_i \rightarrow S_i$ with $\tau_i(s_i) = s_i$. Thus, both $\zeta_i(a, \tau_i(s_i), s_{-i}) = \zeta_i(a, s)$ and $\xi_i(a, \tau_i(s_i), s_{-i}) = \xi_i(a, s)$.¹⁰ Let $\Theta_i = R_i \times A_i$ be the space of pure *deviations* for i .

For every player i and every deviation $\theta_i = (\rho_i, b_i) \in \Theta_i$, the conditional probability of *reported* signals when everyone else is honest and plays $a_{-i} \in A_{-i}$ is given by

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

When all other players are honest and obedient, the utility to i from deviating to $\theta_i = (\rho_i, b_i)$ conditional on being recommended to play a_i under contract (σ, ξ) equals

$$\frac{1}{\sigma(a_i)} \left[\sum_{a_{-i}} v_i(b_i, a_{-i}) \sigma(a) + \sum_{(a_{-i}, s)} \xi_i(a, s) \Pr(s|\theta_i, a_{-i}) \right],$$

where $\sigma(a_i) = \sum_{a_{-i}} \sigma(a) > 0$ is the probability that a_i was recommended. After observing the recommendation to play a_i , player i calculates his conditional expected utility from playing b_i and reporting according to ρ_i by adding payoffs across all possible actions of his opponents and signal realizations given that others will obey the mediator's recommendations and report their monitoring signals truthfully.

¹⁰We will often use the notation $s = (s_i, s_{-i})$ and $a = (a_i, a_{-i})$ for any i , where $s_i \in S_i$ and $s_{-i} \in S_{-i} = \prod_{j \neq i} S_j$; similarly for A_{-i} .

3.1 Linear Metering Problem

A team's *metering problem* is to choose a contract (σ, ξ) consisting of a correlated strategy σ (i.e., with $\sigma \geq 0$ and $\sum_a \sigma(a) = 1$) and a probability weighted incentive scheme ξ that make incentive compatible obeying recommended behavior as well as honest reporting of monitoring signals. For such a contract (σ, ξ) , incentive compatibility is captured by the following family of constraints.

$$\forall i \in I, a_i \in A_i, \theta_i = (\rho_i, b_i) \in \Theta_i, \\ \sum_{a_{-i}} \sigma(a) (v_i(b_i, a_{-i}) - v_i(a)) \leq \sum_{(a_{-i}, s)} \xi_i(a, s) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) \quad (*)$$

The left-hand side reflects the *deviation gain* in terms of utility¹¹ for a player i from playing b_i when asked to play a_i . The right-hand side reflects Mr. i 's *contractual loss* from deviating to $\theta_i = (\rho_i, b_i)$ relative to honesty and obedience (i.e., playing a_i when told to do so and reporting according to τ_i). Such a loss originates from two sources. On the one hand, playing b_i instead of a_i may change conditional probabilities over signals. On the other, reporting according to ρ_i may affect conditional payments.

Definition 3.1. A correlated strategy σ is *approximately implementable* if there exist probability weighted transfers $\xi : I \times A \times S \rightarrow \mathbb{R}$ to satisfy $(*)$ for every (i, a_i, θ_i) .¹²

Next, we will formulate and answer the question: what monitoring technologies allow a given team to overcome its incentive constraints? Naturally, a team's ability to overcome incentives ought to be intimately related to the actions that its monitoring technology is able to distinguish. Thus, if working cannot be distinguished from shirking, there is no hope of finding contracts that motivate someone to work.

Below, we introduce two complementary conditions on a monitoring technology. The first one describes when unilateral deviations can be distinguished. It is shown to be necessary and sufficient for any team with any profile of utility functions to be able to approximately implement any action profile with secret transfers. The second condition describes when obedient players can be identified, and is equivalent to the previous statement with the added restriction of ex post budget balance. Restricted to public monitoring, identifying obedient players is equivalent to the first condition and the existence of transfers that attain any budget without disrupting incentives.

¹¹Specifically, in terms of probability weighted utility, weighted by $\sigma(a_i)$, the probability that a_i was recommended. If a_i is never recommended then $\sigma(a_i) = 0$ and the left-hand side equals zero.

¹²This definition is defended as the limit of exactly implementable strategies in Section 4.1.

3.2 Unilateral Deviations

Definition 3.2. A monitoring technology \Pr *distinguishes unilateral deviations* (DUD) if every vector $\lambda \geq 0$ such that

$$\forall (i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0,$$

must also have the property that for all $\theta_i = (\rho_i, b_i)$, if $\lambda_i(a_i, \theta_i) > 0$ then $b_i = a_i$.

Intuitively, DUD means that any pure or mixed deviation is statistically distinguishable from obedient behavior. To illustrate, fix any player i and let $\Pr[\theta_i] \in \mathbb{R}^{A_{-i} \times S}$ be the vector defined by $\Pr[\theta_i](a_{-i}, s) := \Pr(s|\theta_i, a_{-i})$. For any recommendation-contingent deviation plan $\alpha_i : A_i \rightarrow \Delta(\Theta_i)$, the equation

$$\Pr[\tau_i, a_i] = \sum_{\theta_i \in \Theta_i} \alpha_i(\theta_i|a_i) \Pr[\theta_i]$$

can only be satisfied when the marginal of $\alpha_i(\cdot|a_i)$ on A_i places all its mass on the pure strategy a_i .¹³ Any other α would violate this equality, implying that the statistical distribution of reported signals would be different from what would have arisen had player i played (τ_i, a_i) . This is interpreted as the deviation α_i by player i being statistically detectable when a_i was recommended. In this sense, we allow for dishonesty in reporting strategies to remain undetectable. Such deviations need not be detected as they are not directly profitable.

To compare DUD with other well-known conditions, suppose that monitoring is public (i.e., S_i is a singleton for all $i \neq 0$). In this case, DUD reduces to *convex independence*:

$$\Pr[a_i] \notin \text{conv}\{\Pr[b_i] : b_i \neq a_i\}$$

where $\Pr[b_i] \in \mathbb{R}^{A_{-i} \times S}$ is given by $\Pr[b_i](a_{-i}, s) = \Pr(s|b_i, a_{-i})$. This version of DUD is substantially weaker than the standard condition below:¹⁴

$$\Pr[a] \notin \text{conv}\{\Pr[b_i, a_{-i}] : b_i \neq a_i\},$$

where $\Pr[a] \in \mathbb{R}^S$ is defined by $\Pr[a](s) = \Pr(s|a)$. As is well known, this condition, call it *exact convex independence* (ECI), is necessary and sufficient for exact

¹³To see this, divide both sides of the equation defining DUD by $\sum_{\hat{\theta}_i} \lambda_i(a_i, \hat{\theta}_i)$ for all (i, a, s) .

¹⁴See, for example, Legros and Matsushima (1991), Fudenberg et al. (1994), Compte (1998), Kandori and Matsushima (1998), or Obara (2006a). A popular name for a version of this standard condition has been “individual full rank,” introduced by Fudenberg et al. (1994).

implementation of $a \in A$ (without budget-balance). Intuitively, ECI means that the conditional probability over signals at a given action profile differs from that after any pure or mixed unilateral deviation.

ECI differs from DUD substantially, since DUD is necessary and sufficient for *approximate* implementation (without budget-balance), as will be shown momentarily. Therefore, DUD should constitute a much weaker requirement than ECI. Formally, both ECI and DUD with public monitoring require that certain vectors lie outside the convex hull of some other vectors. For DUD, the vectors have dimension $A_{-i} \times S$, whereas for ECI the vectors are only of dimension S . To illustrate, consider the following example of a (public) monitoring technology satisfying DUD but not ECI. In fact, it violates even *local individual full rank*, introduced by d'Aspremont and Gerard-Varet (1998), which requires individual full rank at some mixed strategy profile, possibly different for each i .

Example 3.3. There are two publicly verifiable signals, $S = \{x, y\}$. There are two players, $I = \{1, 2\}$. Player 1 has two actions, $A_1 = \{U, D\}$, and player 2 has three actions, $A_2 = \{L, M, R\}$. The conditional probability system is given below.

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

In this example, ECI fails, since $\Pr[U, R]$ clearly lies in the convex hull of $\Pr[U, L]$ and $\Pr[U, M]$. Intuitively, there is a mixed deviation (namely $\frac{1}{2}[L] + \frac{1}{2}[M]$, where $[\cdot]$ stands for Dirac measure) by player 2 such that the conditional probability over signals is indistinguishable from what it would be if he played R . In fact, a similar phenomenon takes place when player 1 plays D (this time with mixed deviation $\frac{2}{3}[L] + \frac{1}{3}[M]$) or indeed regardless of player 1's mixed strategy. It is therefore impossible to implement R with standard contracts if player 2 strictly prefers playing L and M , since there always exists a profitable deviation without any contractual losses.

However, it is possible to implement R with secret contracts that correlate player 2's payment with player 1's (recommended) mixed strategy. This way, player 2 will not know with what proportion he ought to mix between L and M in order for his contractual payment to equal what he would obtain by playing R . This suggests how secret, recommendation-contingent rewards can facilitate the most efficient use of a monitoring technology to provide team members with appropriate incentives.

Next, we characterize DUD in terms of approximate implementability.

Definition 3.4. A monitoring technology \Pr *provides strict incentives* (PSI) if given $D_i : A_i \times A_i \rightarrow \mathbb{R}_+$ for all i with $D_i(a_i, a_i) = 0$, there is $\xi : I \times A \times S \rightarrow \mathbb{R}$ such that

$$\forall(i, a_i, \theta_i), \quad D_i(a_i, b_i) \leq \sum_{(a_{-i}, s)} \xi_i(a, s)(\Pr(s|a) - \Pr(s|\rho_i, b_i, a_{-i})).$$

If $D_i(a_i, b_i)$ is interpreted as a player’s *deviation gain* from playing b_i when recommended to play a_i , then PSI means that for any possible deviation gains by the players, there is a contract such that any deviator’s expected contractual loss outweighs his deviation gain after every recommendation. It may appear that PSI is a rather strong condition on a monitoring technology, in contrast with the argued weakness of DUD (see Example 3.3). As the next result shows, it turns out that both conditions are equivalent, in fact mutually dual. (All proofs appear in Appendix B.)

Proposition 3.5. *A monitoring technology distinguishes unilateral deviations if and only if it provides strict incentives.*

This result describes the duality between identifiability and enforceability via secret contracts. The next result, which may be viewed as a corollary, characterizes DUD as the weakest identifiability required for any action to be approximately implementable.

Theorem 3.6. *A monitoring technology distinguishes unilateral deviations if and only if any team with any profile of utility functions can approximately implement any action profile with secret contracts.*

The proof of Theorem 3.6 uses two mutually dual linear programming problems. The *primal* chooses a contract—an allocation— (σ, ξ) to maximize a linear functional f of σ subject to incentive compatibility.¹⁵ The *dual* has contracts as multipliers and motivates Definition 3.2 from a “backward-engineering” exercise: what minimal requirement on a monitoring technology implies that the multipliers λ on incentive constraints equal zero (i.e., incentive constraints do not bind)?

Next, we extend this “backward-engineering” exercise to incorporate budget balance.

¹⁵Although no budget constraints were imposed, we could have imposed *expected* budget balance,

$$\sum_{(i, a, s)} \xi_i(a, s) = 0,$$

but this constraint would not bind, since adding a constant to any ξ preserves its incentive properties.

3.3 Obedient Players

Definition 3.7. A monitoring technology \Pr *identifies obedient players* (IOP) if whenever there exists $\lambda \geq 0$ and $\eta \in \mathbb{R}^{A \times S}$ such that

$$\forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = \eta(a, s),$$

it follows that for all $\theta_i = (\rho_i, b_i)$, if $\lambda_i(a_i, \theta_i) > 0$ then $a_i = b_i$.

IOP is a stronger requirement on a monitoring technology than DUD. Indeed, DUD follows by replacing η above with the zero vector.

What does IOP mean? Considering its defining equation, let

$$\alpha_i(\theta_i|a_i) = \frac{\lambda_i(a_i, \theta_i)}{\sum_{\hat{\theta}_i} \lambda_i(a_i, \hat{\theta}_i)}.$$

Interpret α as the probability that i plays $\theta_i = (\rho_i, b_i)$ when recommended to play a_i . IOP means that any profile of such unilateral deviations that affects the probability of reported signals must do so in a way that is different for different players, since η does not depend on i (unless it comes from dishonest but strictly obedient behavior). Conversely, if IOP fails then there exist disobedient strategies that change conditional probabilities in the same way for every player, so anyone could have been the deviator. Budget-balanced implementation must therefore fail, since players' incentives would "overlap." In other words, it would be impossible to punish some and reward others at the same time in order to provide adequate incentives. If all players must be punished and/or rewarded together, then budget balance must fail.

In comparison with Holmström (1982), who appointed a principal to play the role of budget-breaker, in this model a team whose monitoring technology exhibits IOP can share that role internally. In some teams, this might be allocated stochastically, even leading to a secret principal (Section 2.2).

In the context of *public monitoring*, IOP reduces to DUD and the requirement that

$$\bigcap_{i \in I} C_i = \mathbf{0},$$

where for every i , C_i (called the *cone* of player i) is the set of all vectors $\eta \in \mathbb{R}^{A \times S}$ such that for some $\lambda \geq 0$,

$$\forall(i, a, s), \quad \eta(a, s) = \sum_{b_i \in A_i} \lambda_i(a_i, b_i) (\Pr(s|a) - \Pr(s|b_i, a_{-i})).$$

Call this second condition *non-overlapping cones* (NOC). Fudenberg et al. (1994) impose a full rank condition for each pair of players at each action profile, implying that certain hyperplanes intersect only at the origin for *every pair* of players. On the other hand, NOC requires that certain cones intersect only at the origin for *all* players. So it is possible that two players' cones overlap, i.e., their intersection is larger than just the origin, although we do require that there exist at least two players whose cones do not overlap. Thus, upon a particular unilateral deviation that changes probabilities by DUD, although it may be impossible to identify a single player as the deviator, there must exist at least one player who could not have generated the given statistical change. In this sense, IOP identifies obedient players.¹⁶

Just as with DUD, IOP can be translated to an equivalent condition with dual economic interpretation. This condition takes the form of PSI with budget balance, and its equivalence to IOP follows by the same argument as for DUD and PSI.

Specifically for public monitoring, the fact that IOP can be decomposed into two separate conditions, DUD and NOC, provides useful insights. For instance, we have the following definition and result.

Definition 3.8. A public monitoring technology \Pr *clears every budget* (CEB) if given $K : A \times S \rightarrow \mathbb{R}$ there exists $\xi : I \times A \times S \rightarrow \mathbb{R}$ such that

$$\begin{aligned} \forall(i, a_i, b_i), \quad 0 &\leq \sum_{(a_{-i}, s)} \xi_i(a, s)(\Pr(s|a) - \Pr(s|b_i, a_{-i})), \quad \text{and} \\ \forall(a, s), \quad \sum_{i \in I} \xi_i(a, s) &= K(a, s). \end{aligned}$$

The function $K(a, s)$ may be regarded as a budgetary surplus or deficit for each combination of recommended action and realized signal. CEB means that any level of such budgetary surplus or deficit can be attained by a team without disrupting any incentive compatibility constraints. As it turns out, this is equivalent to NOC.

Proposition 3.9. *A public monitoring technology has non-overlapping cones if and only if it clears every budget.*

This result further clarifies the relative roles of DUD and NOC. By Theorem 3.6, DUD is necessary and sufficient for approximate enforceability of any action profile.

¹⁶IOP is also weaker than pairwise full rank in the sense of approximate versus exact implementation, much like DUD versus individual full rank.

However, the team’s budget may not be balanced ex post (it can only be balanced in expectation). NOC then guarantees existence of a further contract to absorb any budgetary deficit or surplus without disrupting incentive constraints. Therefore, the original contract plus this further contract can implement the same action profile with ex post budget-balance.¹⁷

With private monitoring, a decomposition of IOP into two separate parts does not emerge naturally. This is partly due to the fact that with private monitoring there exist deviations, namely dishonest ones, that do not intrinsically affect anyone’s welfare, and as such we allow them, whereas with public monitoring every deviation may in principle affect players directly.

We now turn to the second main result of this paper, which characterizes approximate implementability with budget balance by IOP. Of course, budget balance means that

$$\forall(a, s), \quad \sum_{i \in I} \xi_i(a, s) = 0. \quad (**)$$

Theorem 3.10. *A monitoring technology identifies obedient players if and only if any team with any profile of utility functions can approximately implement any action profile with budget balanced secret contracts.*

The proof of this result is almost identical to that of Theorem 3.6, therefore omitted. The only difference is that the primal now includes budget balance, which leads to a slightly different dual. Let us now consider some examples to better understand IOP.

Example 3.11. Suppose there exists an individual i_0 such that A_{i_0} and S_{i_0} are both singleton sets. The dual constraints associated with player i_0 are given by

$$\lambda_{i_0}(a_{i_0}, a_{i_0})(\Pr(s|a) - \Pr(s|a)) = \eta(a, s) = 0.$$

It follows that any feasible dual solution must satisfy $\eta(a, s) = 0$ for every (a, s) . Hence, DUD suffices for approximate implementability with ex post budget balance for this team. Since player i_0 cannot be a deviator, she may become a “principal” and serve as “budget-breaker,” much like a seller in an auction.

Example 3.12. Consider a team with two players ($I = \{1, 2\}$) and two publicly verifiable signals ($S = S_0 = \{x, y\}$). The players play the normal-form game (left) with public monitoring technology (right) below:

¹⁷A similar argument is provided by d’Aspremont et al. (2004) for Bayesian mechanisms.

	w	s_2
m	2, -1	-1, 0
s_1	3, -1	0, 0

Utility Payoffs

	w	s_2
m	$p, 1 - p$	$q, 1 - q$
s_1	1/2, 1/2	1/2, 1/2

Signal Probabilities

Suppose that $q > p > 1/2$. First we will show that the “desirable” profile (s_1, w) cannot even be implemented approximately with standard (i.e., non-secret) contracts. With any standard contract, player 1 must be indifferent between monitoring and shirking to approximate efficiency (it can be shown that player 2’s randomization does not help). This implies that $1 = \frac{1}{4}(\zeta_1(x) - \zeta_1(y))$, where $\zeta_1(\omega)$ is the transfer to player 1 when $\omega \in \{x, y\}$ realizes. Budget balance requires $1 = \frac{1}{4}(\zeta_2(y) - \zeta_2(x))$. Since player 2’s incentive constraint is $1 \leq \sigma \frac{1}{4}(\zeta_2(y) - \zeta_2(x))$, where σ denotes the probability that player 1 plays m , it follows that σ cannot be smaller than 1.

There exist budget-balanced secret contracts that approximately implement (s_1, w) . Indeed, let player 1 play m with any probability $\sigma > 0$ and player 2 play w with probability 1. Let $\zeta : A \times S \rightarrow \mathbb{R}$ denote the vector of monetary transfers to player 1 from player 2, and fix $\zeta(a, s) = 0$ for all (a, s) except (m, w, x) . That is, no money is transferred at all except when (m, w) is recommended and x realizes. The incentive constraints associated with recommending s_1 and s_2 are clearly satisfied. The remaining incentive constraints simplify to:

$$\begin{aligned}
 m : \quad & 1 + \frac{\sigma(m, w)}{\sigma(m)} \left(\frac{1}{2} - p\right) \zeta(m, w, x) \leq 0 \\
 w : \quad & 1 + \frac{\sigma(m, w)}{\sigma(w)} (p - q) \zeta(m, w, x) \leq 0
 \end{aligned}$$

These two inequalities can clearly be satisfied by taking $\zeta(m, w, x)$ large enough. It is not difficult to check that both IOP is satisfied (hence also DUD) in this example.

For other p and q , it can be shown that IOP is satisfied if and only if $p \neq q$ and $(p - 1/2)(q - 1/2) > 0$. Thus, if the public signal is perfectly informative about player 2’s behavior (e.g., $\Pr(x|m, w) = \Pr(y|m, s_2) = 1$ as with Robinson and Friday from Section 2.1), approximate implementability with budget balance fails.

With private monitoring (e.g., $S = S_1 = \{x, y\}$), it can be shown that although IOP fails, the same condition suffices for (s_1, w) to be approximately implementable with budget balance. However, not every action profile is approximately implementable. See Section 4.2 for conditions to approximately enforce a given action profile.

4 Discussion

This section makes four comments. Firstly, it fills an important gap in the interpretation of Theorems 3.6 and 3.10. Secondly, it reconciles our main results with the literature by applying the duality of our model to the case of fixed action profiles and utility functions. Thirdly, environmental complications such as limited liability and individual rationality are examined, where standard results generalize to our setting easily, such as that only total liability matters to a team or that individual rationality is not a binding constraint. We end the section by arguing that DUD and IOP, as well as similar variants, are generic in relatively low dimensional spaces.

4.1 Exact versus Approximate Enforcement

A correlated strategy σ is (exactly) *implementable* if there is a scheme ζ such that

$$\forall i \in I, a_i \in A_i, \theta_i \in \Theta_i, \\ \sum_{a_{-i}} \sigma(a)(v_i(b_i, a_{-i}) - v_i(a)) \leq \sum_{(a_{-i}, s)} \sigma(a)\zeta_i(a, s)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})). \quad (***)$$

In Section 3.1, approximate implementability is defined in terms of linear inequalities: σ is approximately implementable if a ξ exists such that (σ, ξ) satisfies (*). To justify, it must be shown that (σ, ξ) is *approachable*: there is a sequence $\{(\sigma^m, \zeta^m)\}$ such that (σ^m, ζ^m) satisfies (***) for every m , $\sigma^m \rightarrow \sigma$, and $\sigma^m \zeta^m \rightarrow \xi$. The next result proves this under DUD and IOP. In addition, IOP implies every action profile is approachable with contracts that are budget balanced “along the way,” not just asymptotically.

Proposition 4.1. *Pr satisfies DUD (IOP) only if every completely mixed correlated strategy is implementable (with budget balance). Hence, DUD (IOP) implies that every contract satisfying (*) (and (**)) is approachable (with budget balance).*

When DUD or IOP fails, the “closure” of (***) does not necessarily equal (*). To illustrate, consider the following variation of Robinson and Friday (Section 2.1):

	work	shirk	rest		work	shirk	rest
monitor	2, -1	-1, 0	-1, 0	monitor	1, 0	0, 1	1, 0
shirk	3, -1	0, 0	0, -1	shirk	1/2, 1/2	1/2, 1/2	1/2, 1/2
	Utility Payoffs				Signal Probabilities		

Assume the signal is public. The profile (shirk,work) is approximately implementable with transfers ξ given by $\xi_F(g|\text{monitor},\text{work}) = 1$ and $\xi_i(a, s) = 0$ for other (i, a, s) . However, since rest is indistinguishable from work and rest weakly dominates work, no contract can dissuade Friday from resting. Hence, (shirk,work) is not approachable. Generalizing Proposition 4.1 involves iterated elimination of weakly dominated *indistinguishable* strategies in the spirit of Myerson's (1997) dual reduction; details are left for another paper. (But Theorem 4.3 below provides a partial generalization.)

4.2 Fixed Action Profiles and Utility Functions

A characterization of implementable action profiles also follows. We focus on budget balanced implementation (without proof, since it is just like that of Theorem 3.10); the unbalanced case—being similar—is omitted. Say Pr *identifies obedient players* at $a \in A$ (IOP- a) if whenever there exists $\lambda \geq 0$ and $\eta \in \mathbb{R}^S$ such that

$$\forall(i, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(\theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = \eta(s),$$

it follows that given $\theta_i = (\rho_i, b_i)$, $\lambda_i(\theta_i) > 0$ implies $b_i = a_i$.

Proposition 4.2. *A monitoring technology identifies obedient players at an action profile a if and only if any team with any profile of utility functions can exactly implement a with budget balanced secret contracts.*

With public monitoring, IOP- a can be decomposed into two conditions. The first is *exact convex independence* at a (ECI- a), which means that the requirement for ECI from Section 3.2 holds at a . For the second, let $C_i(a)$ be the *cone* of player i at a , i.e., the set of all vectors $\eta \in \mathbb{R}^S$ such that for some $\lambda \geq 0$,

$$\forall s \in S, \quad \eta(s) = \sum_{b_i \in A_i} \lambda_i(b_i)(\Pr(s|a) - \Pr(s|b_i, a_{-i})).$$

A public monitoring technology Pr has *non-overlapping cones* at a (NOC- a) if

$$\bigcap_{i \in I} C_i(a) = \mathbf{0}.$$

IOP- a is equivalent to ECI- a and NOC- a . It generalizes the famous *pairwise full rank* condition of Fudenberg et al. (1994), and implies (but is not implied by) for $i \neq j$,

$$C_i(a) \cap C_j(a) = \mathbf{0}.$$

Intuitively, i 's and j 's deviations can be statistically distinguished at a . On the other hand, NOC- a allows some pair of players' cones to overlap. Naturally, this is weaker than pairwise full rank at a , indeed necessary and sufficient for implementation of a when combined with ECI- a . Intuitively, NOC- a requires that some player can be identified as probably obedient.¹⁸

It is possible to partially generalize Proposition 4.1 by fixing utility functions. To this end, \Pr is said to *v-distinguish unilateral deviations* (v -DUD) if given $\lambda \geq 0$,

$$\forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0,$$

implies that $\sum_{\theta_i} \lambda_i(a_i, \theta_i)(v_i(b_i, a_{-i}) - v_i(a)) \leq 0$ for every i and every a . Similarly, \Pr *v-identifies obedient players* (v -IOP) if given $\lambda \geq 0$ and $\eta \in \mathbb{R}^{A \times S}$ for which

$$\forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = \eta(a, s),$$

it follows that $\sum_{(i, \theta_i)} \lambda_i(a_i, \theta_i)(v_i(b_i, a_{-i}) - v_i(a)) \leq 0$ for all $a \in A$. The next result follows immediately from the duality of Theorem 3.6 and Proposition 4.1, so its proof is omitted. It could also be extended to describe exact implementability in line with Proposition 4.2 after suitably amending v -DUD/ v -IOP; details are left to the reader.

Theorem 4.3. *A monitoring technology exhibits v-DUD (v-IOP) if and only if any action profile is approximately implementable with (budget balanced) secret contracts. Furthermore, Proposition 4.1 still holds with v-DUD (v-IOP) replacing DUD (IOP).*

4.3 Participation and Liability

In this subsection we will use duality to study teams subject to liquidity constraints. One such constraint is *limited liability*, where an individual's transfers are bounded below. This can be taken into account by adding $\zeta_i(a, s) \geq \ell_i$ or $\xi_i(a, s) \geq \sigma(a)\ell_i$ to the metering problem, where ℓ_i is an exogenous parameter representing player i 's *liability*. Let $\ell = (\ell_1, \dots, \ell_n)$ be the profile of liabilities faced by a team. A team's *total liability* is defined by $\widehat{\ell} = \sum_i \ell_i$. By a simple duality and without restrictions on a team's monitoring technology, we can generalize to our setting Theorem 5 of Legros and Matsushima (1991) and Theorem 4 of Legros and Matthews (1993).

¹⁸Restricted to public monitoring, Proposition 4.2 is equivalent to Proposition 3 in Legros and Matsushima (1991). Similar results appear also in d'Aspremont and Gérard-Varet (1998), but our decomposition and interpretation are new.

Proposition 4.4. *Only total liability affects a team's (approximately) implementable action profiles (with and without budget balance).*

It is possible that the team faces double-sided limited liability, which may be captured by adding a version of the following constraints to the metering problem:

$$\forall(i, a, s), \quad -\sigma(a)\ell_i \leq \xi_i(a, s) \leq \sigma(a)\ell_i,$$

for some $\ell_i \geq 0$. These constraints lead to an alternative, linear way of requiring that ξ be adapted to σ (i.e., $\xi_i(a, s) = 0$ whenever $\sigma(a) = 0$).

Individual rationality is also amenable to our study of incentives. Without budget balance, since players can be paid lump sums to become indifferent between belonging to the team and forsaking it, individual rationality constraints cannot bind. Hence, suppose the team's budget must be balanced ex post. As a normalization, assume that $\sum_i v_i(a) \geq 0$ for all $a \in A$. Participation constraints may be incorporated as:

$$\forall i \in I, \quad \sum_{a \in A} \sigma(a)v_i(a) + \sum_{s \in S} \xi_i(a, s) \Pr(s|a) \geq 0.$$

Proposition 4.5. *Participation is not a binding constraint if $\sum_i v_i(a) \geq 0$ for all a .*

4.4 Genericity

We conclude this section by arguing that IOP (hence also DUD) holds generically, i.e., for an open set of monitoring technologies whose Lebesgue measure equals that of $\Delta(S)^A$. Intuitively, IOP is generic if the team has access to enough signals. Formally, say \Pr has *maximal rank* if for every player i , the family of vectors

$$\{\Pr[a_i, s_i] \in \mathbb{R}^{A_{-i} \times S_{-i}} : (a_i, s_i) \in A_i \times S_i\}$$

is linearly independent, where $\Pr[a_i, s_i]$ is defined as $\Pr[a_i, s_i](a_{-i}, s_{-i}) = \Pr(s|a)$.

Definition 4.6. A monitoring technology \Pr satisfies *strong convex independence* (SCI) if whenever there exists $\lambda_i : A_i \times S_i \times A_i \rightarrow \mathbb{R}_+$ such that

$$\forall(i, a_i, s_i), \quad \Pr[a_i, s_i] = \sum_{(t_i, b_i)} \lambda_i(a_i, t_i, b_i) \Pr[b_i, t_i],$$

it follows that $\lambda_i(a_i, t_i, b_i) > 0$ implies $b_i = a_i$.

The maximal rank condition is clearly stronger than SCI, and it is not difficult to show that SCI implies IOP. Therefore, if having maximal rank is generic then so is IOP. Since the set of full rank matrices is generic in any Euclidean space, having maximal rank is generic when $|A_i \times S_i|$ is less than or equal to $|A_{-i} \times S_{-i}|$ and neither $|S_i|$ nor $|S_{-i}|$ is a singleton for any player i (Obara, 2006b, Theorem 11). In conclusion:

Proposition 4.7. *IOP is generic if for every player i , both $|A_i \times S_i| \leq |A_{-i} \times S_{-i}|$ and $\min\{|S_i|, |S_{-i}|\} \geq 2$.*

5 Extension

Let us now extend the results of Section 3 to environments with infinitely many actions and signals. We begin with a study of continuous games. Using duality as before, we characterize the set of correlated equilibria and correlated equilibrium payoffs as subdifferentials of some value function (Theorem 5.7). This result is used later in Section 5.3. Finally, we extend the metering problem to monitoring technologies with infinitely many signals. Appendix A contains preliminary results and notation.

5.1 Equilibrium with Linear Programming

The set $I = \{1, \dots, n\}$ of individuals is still finite. Let A_i be a compact Hausdorff space of individual actions for every $i \in I$, and

$$A = \prod_{i=1}^n A_i$$

be the product space of action profiles, endowed with the product topology (also compact Hausdorff). Every individual i has a continuous utility function $v_i : A \rightarrow \mathbb{R}$. A *game* is any triple (I, A, v) as above.

Definition 5.1. A *correlated equilibrium* for the game (I, A, v) is any $\sigma \in \Delta(A)$ such that for every $i \in I$ and every Borel measurable deviation $\beta_i : A_i \rightarrow A_i$,

$$\int_A v_i(\beta_i(a_i), a_{-i}) - v_i(a) \sigma(da) \leq 0.$$

This definition follows Hart and Schmeidler (1989). An alternative description of equilibrium is obtained by restricting attention to continuous deviations, as claimed by the next lemma. It is proved by Hart and Schmeidler (1989, p. 24).

Lemma 5.2. *A correlated strategy $\sigma \in \Delta(A)$ satisfies the incentive constraints*

$$\int_A v_i(\beta_i(a_i), a_{-i}) - v_i(a) \sigma(da) \leq 0$$

for all i and all continuous $\beta_i : A_i \rightarrow A_i$ if and only if σ is a correlated equilibrium.

In order to fit the study of correlated equilibrium into a linear programming framework, we will “immerse” the space of pure-strategy deviations into that of mixed-strategy ones, in the spirit of Hart and Schmeidler’s (1989) proof of their Theorem 3. In light of Lemma 5.2, we would like that every continuous pure-strategy deviation remain continuous when immersed. The next lemma suggests that we may immerse pure-strategy deviations in the space $C(A_i, M(A_i))$ of all weak* continuous functions $\lambda_i : A_i \rightarrow M(A_i)$, i.e., such that $a_i^m \rightarrow a_i$ implies that for any continuous $g : A_i \rightarrow \mathbb{R}$,

$$\int_{A_i} g(b_i) \lambda_i(db_i|a_i^m) \rightarrow \int_{A_i} g(b_i) \lambda_i(db_i|a_i).$$

Lemma 5.3. *If $\beta_i : A_i \rightarrow A_i$ is continuous then $\lambda_i : A_i \rightarrow M(A_i)$ defined by*

$$\forall a_i \in A_i, \quad \lambda_i(a_i) = [\beta_i(a_i)],$$

belongs to $C(A_i, M(A_i))$, where $[b_i]$ stands for Dirac measure at $b_i \in A_i$.

The next result uses this immersion of deviations to describe correlated equilibrium with linear inequalities. Let $C(A_i, M(A_i))_+$ be the positive cone of $C(A_i, M(A_i))$, i.e., the set of $\lambda_i \in C(A_i, M(A_i))$ such that $\lambda_i(a_i)$ is a positive measure for all $a_i \in A_i$.

Proposition 5.4. *A correlated strategy $\sigma \in \Delta(A)$ satisfies the linear inequalities*

$$\int_A \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i) \sigma(da) \leq 0$$

for all $i \in I$ and all $\lambda_i \in C(A_i, M(A_i))_+$ if and only if it is a correlated equilibrium.

$C(A_i, M(A_i))$ is an ordered, locally convex space (Appendix A defines its topology) with a dual, denoted by $C(A_i, M(A_i))^*$. The next result associates the incentive constraints above with a continuous linear operator that maps measures on A to continuous linear functionals on $C(A_i, M(A_i))$, i.e., elements of $C(A_i, M(A_i))^*$.

Lemma 5.5. *For every i , a continuous linear operator $F_i^* : M(A) \rightarrow C(A_i, M(A_i))^*$ is defined by the family of evaluations indexed by $\sigma \in M(A)$ and $\lambda_i \in C(A_i, M(A_i))$*

$$F_i^*(\sigma)(\lambda_i) = \int_A \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i) \sigma(da)$$

when both $M(A)$ and $C(A_i, M(A_i))^$ are endowed with their weak* topologies.*

Therefore, the incentive constraints of Proposition 5.4 may be written as

$$F^*(\sigma) \leq 0 \quad \text{in} \quad \prod_{i=1}^n C(A_i, M(A_i))^*,$$

where $F^* = (F_1^*, \dots, F_n^*)$ and for every player i , F_i^* is the operator of Lemma 5.5. This notation is meant to suggest that correlated equilibrium is the solution of a linear program. Indeed, it solves the following problem, called the *dual*.

$$\sup_{\sigma \geq 0} \int_A f(a) \sigma(da) \quad \text{s.t.} \quad \sigma \in M(A), \quad \sigma(A) = 1, \quad \text{and} \quad F(\sigma) \leq 0 \quad \text{in} \quad \prod_{i=1}^n C(A_i, M(A_i))^*.$$

The objective function $f : A \rightarrow \mathbb{R}$ is assumed to be continuous. Therefore, it is a (weak*) continuous linear functional on $M(A)$. This linear program picks σ subject to it being a probability measure and satisfying incentive compatibility.

Let U be the value function of the following linear program, called the *primal*.

$$U(f) := \inf_{\lambda \geq 0, \kappa} \kappa \quad \text{s.t.} \quad \lambda \in \prod_{i=1}^n C(A_i, M(A_i)), \quad \kappa \in \mathbb{R},$$

$$\forall a \in A, \quad \kappa + \sum_{i=1}^n \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i) \geq f(a) \quad \text{in} \quad C(A).$$

It is not hard to see that this primal is well defined. Indeed, it suffices to show that for every λ , (i) $\sum_i \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i)$ is a continuous real-valued function of a and (ii) $\lambda^m \rightarrow \lambda$ in $\prod_i C(A_i, M(A_i))$ and $\kappa^m \rightarrow \kappa$ implies $F(\lambda^m, \kappa^m) \rightarrow F(\lambda, \kappa)$, where $F(\lambda, \kappa) = \kappa + \sum_{i=1}^n \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i)$. Like Lemma 5.5, this follows because v_i is continuous for every i , hence F is a continuous linear operator whose range is contained in the space of continuous functions on A . It is also not difficult to see that the dual of this primal is the purported dual above.

The value function U has the space of continuous functions as its domain. Since correlated equilibria exist by Theorem 3 of Hart and Schmeidler (1989), it follows that the dual is bounded below. Therefore, U is bounded, since the value of the primal always exceeds that of the dual. Clearly, U is weakly increasing, i.e., $f \leq g$ in $C(A)$ implies that $U(f) \leq U(g)$. In fact, it is Lipschitz, so has bounded steepness, as the next result shows.

Lemma 5.6. *The value function U is Lipschitz: $|U(f) - U(g)| \leq \|f - g\|$.*

By Theorem A.1, U is subdifferentiable, and not only is there an optimum solution to the dual, but also there is no duality gap. The next result now follows.

Theorem 5.7. *The set of correlated equilibria of (I, A, v) is the subdifferential of U evaluated at the zero function. The set of correlated equilibrium payoffs is given by the subdifferential of V at 0, denoted by $\partial V(0)$, where*

$$V(w) := U \left(\sum_{i=1}^n w_i v_i \right).$$

Since $V(0) = 0$, it follows that $\partial V(0) = \{\pi \in \mathbb{R}^n : \forall w \in \mathbb{R}^n, \pi \cdot w \leq V(w)\}$.

This result shows how linear programming is useful in the study of correlated equilibria. The subdifferential characterization of equilibria and equilibrium payoffs is a direct application of the no-gap result. Incidentally, the proof of Lemma 5.6 is reminiscent (a dual version) of the proof of Proposition 3 in Gretsky et al. (1999). As a final remark, we could have proved existence of correlated equilibrium with linear programming instead of relying on Hart and Schmeidler (1989) by describing ε -equilibria with linear inequalities and choosing ε to minimize some linear objective, but the proof would have added nothing substantial to the original.

5.2 Continuous Metering Problem

Next, we will attack a general version of the metering problem introduced previously. Guided by duality, we restrict attention to “well-behaved” monitoring technologies.

A *monitoring technology* is a pair (S, Pr) such that S is the product

$$S = \prod_{j=0}^n S_j$$

of measurable spaces and $\text{Pr} : A \rightarrow \Delta(S)$. A *reporting strategy* for any player i is a measurable map $\rho_i : S_i \rightarrow S_i$. The space of reporting strategies is denoted by R_i . As usual, write $\Theta_i := R_i \times A_i$ with typical element $\theta_i = (\rho_i, b_i)$.

If all other players report honestly and play a_{-i} but player i deviates to $\theta_i = (\rho_i, b_i)$ then the probability that the final report belongs to the Borel set $T \subset S$ is given by

$$\text{Pr}(T|\theta_i, a_{-i}) = \int_S \mathbf{1}_T(s_0, s_1, \dots, \rho_i(s_i), \dots, s_n) \text{Pr}(ds|a).$$

Assumption 5.8. Every S_j is a compact Hausdorff space endowed with the Borel σ -algebra, and S is endowed with the Borel σ -algebra arising from its product topology. Pr is weak* continuous, i.e., $\text{Pr} \in C(A, M(S))$.

By Tychonoff's Theorem, the set of all maps from S_i to itself is compact in the product topology (i.e., the topology of pointwise convergence), and the subspace R_i of Borel measurable maps, being closed, is also compact Hausdorff. It follows that Θ_i is a compact Hausdorff space and $\Pr \in C(\Theta_i \times A_{-i}, M(S))$ for every player i .

Definition 5.9. A monitoring technology \Pr *reconciles infinitesimal deviations* (RID) if for every player i , recommendation $a_i \in A_i$, and deviation $\alpha_i \in \Delta(\Theta_i)$,

$$\sup_{a_{-i}} \left\| \int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \alpha_i(d\theta_i) \right\| \geq c_i$$

for some $c_i \in (0, \infty)$ whenever the left-hand side is nonzero.

Fix any player i and recommendation a_i . RID requires that for any deviation α_i by player i , the monitoring technology must have the largest norm (with respect to others' behavior) of the difference in probability measures over signals between obeying and disobeying bounded below by some constant. RID is a condition on a team's monitoring technology guaranteeing that if signal probabilities change by a deviation then they will change by a minimally significant amount. Intuitively, it ensures that infinitesimal deviations are detectable whenever "non-infinitesimal" ones are, as Step 2b in the proof of Theorem 5.11 below suggests.

RID is a relatively weak condition. For instance, it is satisfied by the monitoring technology of Legros and Matthews (1993), based on deterministic, publicly verifiable output. Thus, $c_i = 2$ in their case. More generally, RID is satisfied by technologies similar to Legros and Matthews (1993) with stochastic rather than deterministic output as well as ones that are not public.¹⁹

A *contract* is a pair (σ, ξ) with $\sigma \in \Delta(A)$ a correlated strategy and $\xi \in M(A, C(S))^I$ a profile of secret (probability-weighted) transfers. The *metering problem* is to find a contract subject to incentive compatibility according to the linear inequalities below.

$$\forall i \in I, \lambda_i \in C(A_i, M(\Theta_i))_+, \quad \int_A \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i | a_i) \sigma(da) \leq \int_{A \times S} \int_{\Theta_i} \Pr(ds | a) - \Pr(ds | \theta_i, a_{-i}) \lambda_i(d\theta_i | a_i) \xi_i(s, da). \quad (\dagger)$$

Like (*), intuitively (\dagger) requires of a contract (σ, ξ) that a player's deviation gain be outweighed by its contractual loss for any possible deviation.

¹⁹For instance, consider a public monitoring technology such that $\Pr(a)$ has a density with a jump discontinuity at some s that "moves" with a . For private monitoring technologies we require that where some of the discontinuities cannot be smoothed out by an individual's report.

Definition 5.10. A monitoring technology \Pr *distinguishes unilateral deviations* (DUD) if for every player i and $\lambda_i \in C(A_i, M(\Theta_i))_+$,

$$\sup_{a_{-i}} \left\| \int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \lambda_i(d\theta_i|a_i) \right\| = 0$$

only if $\text{supp } \lambda_i(a_i) \subset R_i \times \{a_i\}$ for every i and every a_i .

This is the continuum analogue of DUD in the finite case, with similar result below.

Theorem 5.11. *Suppose a monitoring technology reconciles infinitesimal deviations. It also distinguishes unilateral deviations if and only if any team with any profile of continuous utility functions can approximately implement any action profile with secret transfers.*

We conclude this section with the version of this theorem subject to budget-balance. It is a straight-forward extension of the above. Of course, budget balance means that

$$\forall i \in I, s \in S, \varphi \in C(A)_+, \quad \sum_{i=1}^n \int_A \varphi(a) \xi_i(s, da) = 0. \quad (\ddagger)$$

Definition 5.12. A monitoring technology \Pr *identifies obedient players* (IOP) if for every $\lambda \in \prod_i C(A_i, M(\Theta_i))_+$ and $\eta \in C(A, M(S))$,

$$\int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \lambda_i(d\theta_i|a_i) = \eta(a)$$

only if $\text{supp } \lambda_i(a_i) \subset R_i \times \{a_i\}$ for every i and every a_i .

Theorem 5.13. *Suppose a monitoring technology reconciles infinitesimal deviations. It also identifies obedient players if and only if any team with any profile of continuous utility functions can approximately implement any action profile with secret, budget-balanced transfers.*

The two results above, with and without budget balance, are significant generalizations of the work of Legros and Matthews (1993) to stochastic, private monitoring technologies. As a final remark, imposing joint conditions on \Pr and v would permit a weakening of RID to establish the characterization results above.²⁰

²⁰For instance, if v_i is Lipschitz, then intuitively it suffices for the steepness of \Pr to be bounded below by the steepness of v_i .

5.3 Repeated Games with Private Monitoring

Consider the problem of repeated interaction with mediated communication. Without explicit monetary transfers, incentives are provided only with changes to continuation payoffs as a result of contingent future play. Applying Theorem 5.7, we characterize equilibrium payoffs and show that SCI (Definition 4.6) suffices to enforce approximate efficiency, indeed a folk theorem.

The timing of the game runs as follows. In each period, a mediator makes confidential, non-binding recommendations to players based on past recommendations, public signals, and reports. Players then simultaneously choose actions. After taking actions, players observe private signals and report to the mediator. Public signals realize²¹ and timing repeats itself ad infinitum. Define the following *partial histories*:

Period	Mediator		Player i	
1	$H_1 = \{0\}$	$K_1 = A$	$H_1^i = \{0\}$	$K_1^i = A_i$
2	$H_2 = A \times S$	$K_2 = H_2 \times A$	$H_2^i = A_i \times S_i$	$K_2^i = H_2^i \times S_0 \times A_i$
t	$H_t = K_{t-1} \times S$	$K_t = H_t \times A$	$H_t^i = K_{t-1}^i \times S_i$	$K_t^i = H_t^i \times S_0 \times A_i$

K_t collects the mediator's partial histories up to and including his recommendations in period t , whereas H_{t+1} describes histories up to and including public signal realizations and private signal reports at t . Similarly, K_t^i and H_{t+1}^i describe, respectively for every t , player i 's *private* partial histories up to the previous period's public signal and his current actual action, and up to and including his privately observed signal.

A *pure communication strategy* is any history-contingent recommendation plan for the mediator, $\mathbf{a} : \bigcup_{t \geq 1} H_t \rightarrow A$. Let \mathcal{A} be the set of all such pure communication strategies, and $\Delta(\mathcal{A})$ be the set of (correlated) communication strategies with typical element σ . For every player i , a *pure deviation* is a pair $\theta^i = (\rho^i, b^i)$ such that

$$b^i : \bigcup_{t \geq 1} K_t^i \rightarrow A_i, \quad \text{and} \quad \rho^i : \bigcup_{t \geq 1} H_t^i \rightarrow S_i.$$

Let Θ^i be the set of all pure deviations for player i . We will now make the following simplifying assumptions on a team's monitoring technology. Every period, the conditional probability of observed signals depends only on players' contemporaneous profile of actions. Additionally, in order to ignore sequential rationality constraints, suppose every realization has positive probability.

²¹We could have allowed player i to observe the public signal at the same time as the private one. We chose this signal timing to stay consistent with previous sections.

Assumption 5.14. $\Pr(s|a) > 0$ for all (a, s) .

If every player is honest and obedient and the mediator plays the pure communication strategy \mathbf{a} , we will write as follows the conditional probability of partial histories:

$$\Pr(h_{t+1}|\mathbf{a}) = \Pr(a_1, s_1, \dots, a_t, s_t|\mathbf{a}) = \prod_{1 \leq u \leq t} \Pr(s_u|a_u)$$

if $a_u = \mathbf{a}(h_u)$ for $1 < u \leq t$, and 0 otherwise. Similarly, $\Pr(k_t|\mathbf{a}) = \sum_{s_t} \Pr(h_{t+1}|\mathbf{a})$.

If player i deviates to θ^i but everyone else is honest and obedient, the joint probability of the mediator's partial history $h_{t+1} = (a_1, s_1, \dots, a_t, s_t)$ and player i 's private history $\widehat{h}_{t+1}^i = (\widehat{a}_1^i, \widehat{s}_1^i, \dots, \widehat{a}_t^i, \widehat{s}_t^i)$ is written as

$$\Pr(h_{t+1}, \widehat{h}_{t+1}^i | \theta^i, \mathbf{a}) = \prod_{1 \leq u \leq t} \Pr(\widehat{s}_u^i, s_u^{-i} | \widehat{a}_u^i, a_u^{-i})$$

if $a_u = \mathbf{a}(h_u)$, $s_u^i = \rho_i(\widehat{h}_u^i)$, and $\widehat{a}_u^i = b_i(\widehat{h}_u^i, \mathbf{a}_i(h_u))$; otherwise $\Pr(h_{t+1}, \widehat{h}_{t+1}^i | \theta^i, \mathbf{a}) = 0$. Similarly, $\Pr(k_t, \widehat{k}_t^i | \theta^i, \mathbf{a}) = \sum_{(s_t, \widehat{s}_t^i)} \Pr(h_{t+1}, \widehat{h}_{t+1}^i | \theta^i, \mathbf{a})$.

Definition 5.15. A communication strategy σ is a *communication equilibrium* if everyone prefers to be honest and obedient when all others are, too: $\forall i \in I, \theta^i \in \Theta^i$,

$$\sum_{\mathbf{a} \in \mathcal{A}} \sigma(\mathbf{a}) \sum_{(t, h_t, \widehat{h}_t^i)} \delta^{t-1} [v_i(b^i(\widehat{h}_t^i, \mathbf{a}_i(h_t)), \mathbf{a}_{-i}(h_t)) \Pr(h_t, \widehat{h}_t^i | \theta^i, \mathbf{a}) - v_i(\mathbf{a}(h_t)) \Pr(h_t | \mathbf{a})] \leq 0.$$

A direct application of Theorem 5.7 yields the following result.

Corollary 5.16. *Given $\delta \in [0, 1)$ and $w \in \mathbb{R}^n$, define the following value function:*

$$V_\delta(w) := \inf_{\lambda \geq 0, \kappa} \kappa \text{ s.t. } \forall \mathbf{a} \in \mathcal{A}, \quad \kappa \geq (1 - \delta) \left\{ \sum_{(i, t, h_t)} \delta^{t-1} w_i v_i(\mathbf{a}(h_t)) \Pr(h_t | \mathbf{a}) - \sum_{(i, \theta^i)} \lambda_i(\theta^i) \sum_{(t, h_t, \widehat{h}_t^i)} \delta^{t-1} [v_i(b^i(\widehat{h}_t^i, \mathbf{a}_i(h_t)), \mathbf{a}_{-i}(h_t)) \Pr(h_t, \widehat{h}_t^i | \theta^i, \mathbf{a}) - v_i(\mathbf{a}(h_t)) \Pr(h_t | \mathbf{a})] \right\}.$$

The set of communication equilibrium payoffs of the repeated game above is $\partial V_\delta(0)$, where $\partial V_\delta(0) = \{\pi \in \mathbb{R}^n : \forall w \in \mathbb{R}^n, w \cdot \pi \leq V_\delta(w)\}$ is the subdifferential of V_δ at 0.

The subdifferential $\partial V_\delta(0)$ characterizes communication equilibrium payoffs of discounted repeated games with imperfect private monitoring and private strategies (and signal probabilities with full support). It would be interesting to find a family of *finitely* many inequalities to approximate $\partial V_\delta(0)$; we leave this for another paper.

Next, we establish a *folk theorem* with mediated communication. To this end, we focus on a particular type of equilibrium called *T-communication equilibrium*. In this equilibrium, a repeated game is divided into many *T*-period blocks. The mediator's recommendations do not depend on recommended actions or reported private signals within a block (i.e., recommendations are *autonomous*, but players still report their private signals every period. At the end of a *T*-period block, the mediator reveals recommended actions and reports of private signals publicly, and a new *T*-period block begins ad infinitum.²² *T*-communication equilibrium is recursive every *T* periods. Any continuation equilibrium starting from period $T + 1$ is isomorphic to an equilibrium of the original game. We use this recursive structure to characterize the set of *T*-communication equilibrium payoffs when players become infinitely patient.

We begin defining *T*-period strategies. A map $\mathbf{a}_T : \bigcup_{t=1}^T H_t \rightarrow A$ is called a *T-period pure communication strategy* if it is constant on each H_t for $t = 1, 2, \dots, T$. A *T-period deviation* for player i is denoted by $\theta_T^i = (\rho_T^i, b_T^i)$, and consists of any truncated version of θ^i to at most *T*-period partial histories. An upper bound of the weighted *T*-communication equilibrium payoffs in the direction $w \in \mathbb{R}^n$ is given by the following linear program:

$$\begin{aligned}
V_\delta^T(w) = & \sup_{\sigma \geq 0, \xi} \sum_{\mathbf{a}_T \in \mathcal{A}_T} \sigma(\mathbf{a}_T) \sum_{(i,t,h_t)} \frac{(1-\delta)\delta^{t-1}}{1-\delta^T} w_i v_i(\mathbf{a}_T(h_t)) \Pr(h_t|\mathbf{a}_T) \\
& + \sum_{(i,\mathbf{a}_T,h_{T+1})} w_i \xi_i(h_{T+1}, \mathbf{a}_T) \Pr(h_{T+1}|\mathbf{a}_T) \quad \text{s.t.} \quad \sum_{\mathbf{a}_T} \sigma(\mathbf{a}_T) = 1, \\
& \forall h_{T+1}, \quad \sum_{i=1}^n w_i \xi_i(h_{T+1}, \mathbf{a}_T) \leq 0, \quad \forall (i, \theta_T^i), \\
& \sum_{\mathbf{a}_T} \sigma(\mathbf{a}_T) \sum_{(t,h_t,\hat{h}_t^i)} \delta^{t-1} [v_i(b^i(\hat{h}_t^i, \mathbf{a}_T^i(h_t)), \mathbf{a}_T^{-i}(h_t)) \Pr(h_t, \hat{h}_t^i|\theta^i, \mathbf{a}_T) - v_i(\mathbf{a}_T(h_t)) \Pr(h_t|\mathbf{a}_T)] \\
& \leq \frac{1-\delta^T}{1-\delta} \sum_{(\mathbf{a}_T, h_{T+1})} [\xi_i(h_{T+1}, \mathbf{a}_T) \Pr(h_{T+1}|\mathbf{a}_T) - \xi_i(h_{T+1}, \mathbf{a}_T) \Pr(h_{T+1}|\mathbf{a}_T, \theta_T^i)].
\end{aligned}$$

The problem above can be expressed in the familiar form of Fudenberg et al. (1994) by letting $\xi_i(h_{T+1}, \mathbf{a}_T) = \frac{\delta^T}{1-\delta^T} (\sigma(\mathbf{a}_T) V_i(h_{T+1}, \mathbf{a}_T) - \bar{v}_i)$, where $V_i(h_{T+1}, \mathbf{a}_T)$ stands for player i 's continuation payoff given (h_{T+1}, \mathbf{a}_T) and \bar{v}_i is player i 's total payoff. The recursive constraints $\sum_i w_i \xi_i(h_{T+1}, \mathbf{a}_T) \leq 0$ imply that the weighted sum of

²²A *T*-communication equilibrium would still be an equilibrium even if accumulated information were not made public. If information were not revealed, the mediator would know the continuation equilibrium to be played, but players would not. Their incentive constraints would still be satisfied by virtue of being less informed.

continuation payoffs must be less than or equal to the maximized weighted sum of individual lifetime payoffs. Since this is the only constraint on feasible payoffs besides incentive compatibility, it follows that $V_\delta(w)$ is an upper bound for the weighted sum of T -communication equilibrium payoffs.

By the Maximum Theorem, $\lim_{\delta \rightarrow 1} V_\delta^T(w)$ exists; call it $V_1^T(w)$. Call its subdifferential $\partial V_1^T(0) := \{\pi \in \mathbb{R}^n : \forall w \in \mathbb{R}^n, w \cdot \pi \leq V_1^T(w)\}$, and denote by E_δ^T the set of T -communication equilibrium payoffs.

Lemma 5.17. *If $\partial V_1^T(0)$ is a convex body²³ in \mathbb{R}^n then every smooth compact convex set W such that $W \subset \text{int } \partial V_1^T(0)$ also satisfies*

$$W \subset \liminf_{\delta \rightarrow 1} E_\delta^T.$$

This follows immediately from Fudenberg et al. (1994) by viewing each T -period block as a one-shot game, hence we omit a formal proof.

Below, we establish a folk theorem with communication by showing that $\partial V_1^T(0)$ covers all payoff profiles above any correlated equilibrium payoff profile as T tends to ∞ . To illustrate how this result might be applied, let us prove the folk theorem for the example of Section 2.2 with 1-communication equilibria. We will show that an optimal outcome in any direction w can be supported with budget-balance.²⁴ To approximately enforce the action profile $(1, \dots, 1)$, which is optimal for $w = (1/n, \dots, 1/n)$, use the continuation payoffs from the contracts in Section 2.2. This works for any positive w whose optimal action profile is $(1, \dots, 1)$. If it is optimal to assign different actions across players for positive w , this can be achieved as usual, without the need for secret randomization. Next, suppose w has both positive and negative components. For example, say $w_i > 0$ and $w_j < 0$ and the optimal action is 1 for player i and 0 for player j . In this case, we can use simultaneous punishments as “budget-balanced” transfers. Secret randomization is helpful here: With high probability, the recommended action profile $(a_i, a_j) = (1, 0)$ is optimal and no transfers are made. If $(a_i, a_j) = (1, 1)$ is recommended, impose a mutual punishment upon the realization of bad news. It is easy to check that this scheme is incentive compatible and approximately optimal. When the same actions are assigned to i and j , again secret randomization is not necessary. Finally, we can always achieve the stage game Nash equilibrium with no transfers for negative w . Combining all these cases, it follows that V_1^1 covers all the individually rational and feasible payoffs.

²³I.e., a convex set with nonempty interior.

²⁴Budget-balance implies dynamic efficiency in repeated games.

We will now establish a folk theorem with respect to payoffs from any stage-game correlated equilibrium σ^* . Let $V = \text{conv} \{v(a) : a \in A\}$, where $v(a) = (v_1(a), \dots, v_n(a))$, be the set of feasible payoffs, and denote the feasible payoffs in excess of those from σ^* by $V(\sigma^*) = \{v \in V : \forall i \in I, v_i \geq \sum_a \sigma^*(a)v_i(a)\}$. For SCI, see Definition 4.6.

Theorem 5.18. *If Pr has SCI then given a stage-game correlated equilibrium σ^* ,*

$$V(\sigma^*) \subset \liminf_{T \rightarrow \infty} \liminf_{\delta \rightarrow 1} E_\delta^T$$

whenever $V(\sigma^)$ has full dimension.*

In a recent paper, Tomala (2005) proves a folk theorem (with respect to correlated minmax payoffs) with private monitoring and “semi-public” communication, in that recommendations and reports are publicly revealed at the end of every period. He relaxes the sufficient conditions of Kandori and Matsushima (1998) by allowing for “intra-period” secrets. Like Kandori and Matsushima (1998), he imposes conditions on the monitoring technology for each pair of players. Intuitively he requires that unilateral deviations by any pair of players be distinguishable, i.e., one can identify which of the two have deviated. On the other hand, we only need that an obedient player be identified for any profile of unilateral deviations. In addition, we allow for secrets to drag through time indefinitely. As such SCI yields a stronger folk theorem than Tomala (2005) (who also needs at least three players). Aoyagi (2005) also proves a folk theorem with private monitoring and communication. He assumes that monitoring is jointly ε -perfect, i.e., players’ actions are almost perfectly observed once other players’ signals are aggregated. SCI is much weaker. Finally, the folk theorem established here generalizes Obara (2006a) to allow mediated communication.

6 Conclusion

In this paper we have explored possible ways in which secret contracts may help organizations, with particular emphasis on the question of monitoring a monitor and maintaining budget balance. Formally, we have used duality systematically to make general statements about a team’s contractual scope. We have exploited this duality to consider teams with infinitely many actions and signals, with fruitful applications such as a subdifferential characterization equilibrium payoffs. Below, we conclude this paper with some comments to connect the paper with the (mechanism design and implementation) literature, discuss weaknesses (collusion), and further research.

6.1 Abstract Mechanisms in Concrete Contracts

We build a bridge between abstract mechanism design and concrete contract theory in this paper. Much of the mechanism design literature has focused on surplus extraction in environments with adverse selection. Thus, Cremer and McLean (1988) argued that if individuals have “correlated types” then their surplus may be extracted.²⁵ On the other hand, they do not explain the source of such correlation. Secret contracts provide an explanation for the emergence of correlated types.

As part of a team’s economic organization, it may be beneficial for private information to be allocated differently in order to provide the right incentives. As has been argued here, this is true even if the team starts without informational asymmetry. In a sense, correlated types emerge endogenously, and as such there are incidental similarities between this paper and the mechanism design literature even if conceptually there are important differences. For instance, the essence of secret contracts is lost in the abstraction of mechanism design because it is so reduced. With moral hazard, our identifiability conditions apparently lend themselves easily to interpretation.

Nonetheless, a hybrid exercise where players begin with some private information and face an additional metering problem is amenable to the techniques developed here. Initial results are promising (Rahman, 2005b, Ch. 5); details are for another paper.

6.2 Secrets and Verifiable Recommendations

Secret contracts rely on making payments contingent on verifiable recommendations. However, without a disinterested mediator or machine, this may no longer be a tenable assumption. In this case, it may be possible for players to verifiably reveal their recommendations. Player i ’s reporting strategy would involve announcing a recommended action and a private signal given the true recommendation, his action and his true private signal. Incentive constraints are only slightly different.

Kandori (2003) used similar schemes in repeated games with public monitoring. Our conditions are weaker because we study private monitoring with communication in correlated equilibrium. Moreover, we do not impose pairwise conditions on the monitoring technology as in Kandori (2003) or Tomala (2005) to prove the folk theorem.

²⁵d’Aspremont et al. (2004) extend this result to include budget balance. The additional constraint of individual rationality is studied by Kosenok and Severinov (2004).

As illustrated by Robinson and Friday in Section 2.1, secret contracts provide an intuitive organizational design. If recommendations were not verifiable, then in order to approximate efficiency Friday would need to report whether or not he worked, which taken more generally provides a different answer to the question of monitoring the monitor: bring in two monitors and have them monitor each other.

6.3 Usual Problems with Collusion

A notable weakness of secret contracts is not being collusion-proof. To illustrate, in our leading example (Section 2.1) Robinson and Friday could communicate to break down the incentives that secrets tried to provide. However, this problem is neither inherent to secrets nor widespread to all teams. Example 3.12 describes when Robinson and Friday can approximate efficiency with budget balance, for which they require secrets. There, contracts are naturally robust to collusion, since budget balance implies that Friday's gain is Robinson's loss.

Collusion is a problem for secret contracts inasmuch as it is a problem for contracts in general. For instance, the transfer schemes of Cremer and McLean (1988) are not generally collusion-proof for similar reasons. In any case, although there may be partial solutions to the problem of collusion with secret contracts in the spirit of, say, Che and Kim (2006), the main purpose of this paper is to introduce secret contracts. Thus, an analysis of collusion is postponed for the future.

References

- Alchian, A. and H. Demsetz (1972): "Production, Information Costs, and Economic Organization," *American Economic Review*, 62, 777–795.
- Aoyagi, M. (2005): "Collusion Through Mediated Communication in Repeated Games with Imperfect Private Monitoring," *Economic Theory*, 25, 455–475.
- Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.
- Che, Y.-K. and J. Kim (2006): "Robustly Collusion-Proof Implementation," *Econometrica*, 74, 1063–1107.

- Compte, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66, 597–626.
- Cremer, J. and R. McLean (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1257.
- d’Aspremont, C., J. Cremer, and L.-A. Gérard-Varet (2004): “Balanced Bayesian Mechanisms,” *Journal of Economic Theory*, 115, 385–396.
- d’Aspremont, C. and L.-A. Gérard-Varet (1998): “Linear Inequality Methods to Enforce Partnerships under Uncertainty: An Overview,” *Games and Economic Behavior*, 25, 311–336.
- Fudenberg, D., D. Levine, and E. Maskin (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039.
- Gretsky, N., J. Ostroy, and W. Zame (1999): “Perfect Competition in the Continuous Assignment Model,” *Journal of Economic Theory*, 88, 60–118.
- (2002): “Subdifferentiability and the Duality Gap,” *Positivity*, 6, 1–16.
- Hart, S. and D. Schmeidler (1989): “Existence of Correlated Equilibria,” *Mathematics of Operations Research*, 14, 18–25.
- Holmström, B. (1982): “Moral Hazard in Teams,” *Bell Journal of Economics*, 13, 324–340.
- Kandori, M. (2003): “Randomization, Communication, and Efficiency in Repeated Games with Imperfect Public Monitoring,” *Econometrica*, 71, 345–353.
- Kandori, M. and H. Matsushima (1998): “Private Observation, Communication, and Collusion,” *Econometrica*, 66, 627–652.
- Kosenok, G. and S. Severinov (2004): “Individually Rational, Balanced-Budget Bayesian Mechanisms and the Informed Principal Problem,” Mimeo.
- Legros, P. and H. Matsushima (1991): “Efficiency in Partnerships,” *Journal of Economic Theory*, 55, 296–322.
- Legros, P. and S. Matthews (1993): “Efficient and Nearly Efficient Partnerships,” *Review of Economic Studies*, 60, 599–611.
- Myerson, R. (1997): “Dual Reduction and Elementary Games,” *Games and Economic Behavior*, 21, 183–202.

- Nau, R. F. and K. F. McCardle (1990): “Coherent Behavior in Noncooperative Games,” *Journal of Economic Theory*, 50, 424–444.
- Obara, I. (2006a): “Folk Theorem with Communication,” Mimeo.
- (2006b): “The Full Surplus Extraction Theorem with Hidden Actions,” Mimeo.
- Radner, R., R. Myerson, and E. Maskin (1986): “An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria,” *Review of Economic Studies*, 53, 59–69.
- Rahman, D. (2005a): “Optimum Contracts with Public and Private Monitoring,” Mimeo.
- (2005b): “Team Formation and Organization,” Ph.D. dissertation, UCLA.
- Rockafellar, R. T. (1970): *Convex Analysis*, Princeton, New Jersey: Princeton University Press.
- Strausz, R. (1997): “Delegation of Monitoring in a Principal-Agent Relationship,” *Review of Economic Studies*, 64, 337–357.
- Tomala, T. (2005): “On Subgame-Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring,” Mimeo.

A Preliminaries

Let X and Y be ordered, locally convex topological vector spaces, with dual spaces X^* and Y^* , respectively. Let $F : X \rightarrow Y$ be a continuous linear operator, $g \in Y$, and $h^* \in X^*$. Define the *primal* as the following linear program:²⁶ $\inf_{x \in X} \{h^*(x) : F(x) \geq g, x \geq 0\}$. The *dual* is defined as $\sup_{y^* \in Y^*} \{y^*(g) : F^*(y^*) \leq h^*, y^* \geq 0\}$, where $F^* : Y^* \rightarrow X^*$ is the *adjoint* of F . The *value function* associated with the primal is the map $V : Y \rightarrow \mathbb{R}$ that results by viewing the value of the primal as a function of the right-hand side constraints, g . Formally, $V(g) := \inf_{x \in X} \{h^*(x) : F(x) \geq g, x \geq 0\}$. By definition, for a given triple (F, g, h^*) , there is *no duality gap* if the value of the primal equals the value of the dual.

A convex function $f : X \rightarrow \mathbb{R}$ is called *subdifferentiable* at $x \in X$ if its *subdifferential* at x , denoted and defined by $\partial f(x) := \{x^* \in X^* : \forall y \in X, x^*(x) - f(x) \geq x^*(y) - f(y)\}$, is not empty. The result below is taken from Theorem 1 and Section 5 of Gretsky et al. (2002).

²⁶A linear map between ordered vector spaces $T : X \rightarrow Y$ is *positive* if $T(x) \geq 0$ in Y whenever $x \geq 0$ in X . For two linear maps T and U , $T \geq U$ means that $T - U$ is a positive linear map.

Theorem A.1. Fix a triple (F, g, h^*) as above. Both the dual has a solution and there is no duality gap if and only if V is subdifferentiable at g . A convex function f on a normed space Y is subdifferentiable at $z \in Y$ if and only if it has bounded steepness at z , i.e.,

$$\forall y \in Y, \quad f(z) - f(y) \leq C \|z - y\|.$$

Given a compact Hausdorff space X , denote by $C(X)$ the set of all continuous, real-valued functions on X , endowed with the supremum norm. Let $M(X)$ be the space of Radon measures on X with the weak* topology, and $\Delta(X)$ its subspace of probability measures.

Given two compact Hausdorff spaces X and Y , let $C(X, M(Y))$ denote the space of weak* continuous functions on X with values in $M(Y)$, i.e., functions $\mu : X \rightarrow M(Y)$ such that $x_n \rightarrow x$ in X implies $\int_Y g(y)\mu(dy|x_n) \rightarrow \int_Y g(y)\mu(dy|x)$ for every continuous $g : Y \rightarrow \mathbb{R}$. Clearly, $C(X, M(Y))$ is a linear space. It is ordered as follows. If μ and ν belong to $C(X, M(Y))$ then $\mu \leq \nu$ in $C(X, M(Y))$ if for every $x \in X$, $\mu(x) \leq \nu(x)$ in $M(Y)$, i.e., if $\int_Y g(y)\mu(dy|x) \leq \int_Y g(y)\nu(dy|x)$ for every continuous $g : Y \rightarrow [0, \infty)$. It can also be made into a Hausdorff, locally convex topological vector space by endowing it with the weakest topology that makes continuous addition, scalar multiplication, and each of the family of seminorms $\|\mu\|_g = \sup_{x \in X} \int_Y |g(y)| |\mu| (dy|x)$ indexed by $g \in C(Y)$.

The set of continuous linear functionals on $C(X, M(Y))$ endowed with the weak* topology is denoted by $C(X, M(Y))^*$. In this space, convergence is defined as follows. We have $\alpha_n \rightarrow \alpha$ in $C(X, M(Y))^*$ if $\langle \mu, \alpha_n \rangle \rightarrow \langle \mu, \alpha \rangle$ for every μ in $C(X, M(Y))$. It can be shown that $C(X, M(Y))^*$ is identifiable with $M(X, C(Y))$, so $\langle \mu, \alpha \rangle = \int_{X \times Y} \mu(x, dy)\alpha(y, dx)$.

B Proofs

Proposition 3.5. Consider the following linear program.

$$\sup_{\xi} 0 \text{ s.t. } \forall(i, a_i, \theta_i), \quad D_i(a_i, b_i) \leq \sum_{(a_{-i}, s)} \xi_i(a, s)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})).$$

The dual of this problem is given by

$$\inf_{\lambda \geq 0} - \sum_{(i, a_i, \theta_i)} \lambda_i(a_i, \theta_i) D_i(a_i, b_i) \text{ s.t. } \forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0.$$

If PSI is satisfied then the value of the primal is zero. By FTLP, the dual value must be also zero for every $\{D_i\}$ if PSI is satisfied. This implies that $\lambda_i(a_i, \theta_i) = 0$ for every $\theta_i = (\rho_i, b_i)$ with $b_i \neq a_i$ at any feasible dual solution, i.e., DUD. Conversely, if PSI is not satisfied then there exists $\{D_i\}$ for which the primal feasible set is empty, and the primal value is $-\infty$. At

such $\{D_i\}$ there must exist $\lambda \geq 0$ to satisfy the dual constraint and make the dual objective strictly negative, so $\lambda_i(a_i, \theta_i) > 0$ for some (i, a_i, θ_i) with $b_i \neq a_i$, i.e., DUD fails. \square

Theorem 3.6. Consider the following linear program, called the *primal*.

$$V_f(v) := \sup_{\sigma \geq 0, \xi} \sum_{a \in A} f(a) \sigma(a) \quad \text{s.t.} \quad \sum_{a \in A} \sigma(a) = 1,$$

$$\forall (i, a_i, \theta_i), \sum_{a_{-i}} \sigma(a) (v_i(b_i, a_{-i}) - v_i(a)) \leq \sum_{(a_{-i}, s)} \xi_i(a, s) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})).$$

The *dual* is given below. By FTLP, the value of the dual equals that of the primal.

$$V_f(v) = \inf_{\lambda \geq 0, \kappa} \kappa \quad \text{s.t.}$$

$$\forall a \in A, \quad \kappa \geq f(a) - \sum_{(i, \theta_i)} \lambda_i(a_i, \theta_i) (v_i(b_i, a_{-i}) - v_i(a))$$

$$\forall (i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0$$

We will show that DUD is equivalent to $V_f(v) = \max\{f(a) : a \in A\}$ for all f . If Pr satisfies DUD then by the second family of dual constraints, any feasible $\lambda \neq 0$ must have $\lambda_i(a_i, \theta_i) > 0$ only if $a_i = b_i$. Hence, the first family of dual constraints becomes $\kappa \geq f(a)$ for all a . Minimizing κ subject to them yields $\max\{f(a) : a \in A\}$ for any f and v , proving sufficiency. For necessity, if DUD fails there is $\lambda \geq 0$ with

$$\sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0$$

for all (i, a, s) and $\lambda_j(\hat{a}_j, \hat{\theta}_j) > 0$ for some $(j, \hat{a}_j, \hat{\theta}_j)$ with $\hat{b}_j \neq \hat{a}_j$. Let $f = \mathbf{1}_{\hat{a}_j}$ and choose v as follows. For any a_{-j} , the utility to each player depending on whether or not j plays \hat{a}_j is given by (first is j then anyone else):

a_j	\hat{a}_j
1, 0	0, 2

Given a with $a_j \neq \hat{a}_j$, the first dual constraint becomes $0 + \sum_{\rho_j} \lambda(a_j, \hat{a}_j, \rho_j) \leq \kappa$. This can be made smaller than 1 by multiplying λ by a sufficiently small positive number. At \hat{a}_j , the constraint becomes $1 - \sum_{\theta_j} \lambda_j(\hat{a}_j, \theta_j) \leq \kappa$. Since $\sum \lambda > 0$, there is a feasible dual solution with $\kappa < 1 = \max\{f(a)\}$, as required. \square

Proposition 3.9. Consider the following primal problem: Find a feasible ξ to solve

$$\forall (i, a_i, b_i), 0 \leq \sum_{(a_{-i}, s)} \xi_i(a, s) (\Pr(s|a) - \Pr(s|b_i, a_{-i})), \quad \text{and} \quad \forall (a, s), \sum_{i \in I} \xi_i(a, s) = K(a, s).$$

The dual of this problem is given by

$$\inf_{\lambda \geq 0, \eta} \sum_{(a,s)} \eta(a,s) K(a,s) \quad \text{s.t.} \quad \forall (i, a, s), \quad \sum_{b_i \in A_i} \lambda_i(a_i, b_i) (\Pr(s|a) - \Pr(s|b_i, a_{-i})) = \eta(a, s).$$

If CEB is satisfied, then the value of the primal equals 0 for any $K : A \times S \rightarrow \mathbb{R}$. By FTLP, the value of the dual is also 0 for any $K : A \times S \rightarrow \mathbb{R}$. Therefore, any η satisfying the constraint for some λ must be 0 for all (a, s) , so NOC is satisfied. For necessity, if NOC is satisfied then the value of the dual is always 0 for any $K : A \times S \rightarrow \mathbb{R}$. By FTLP, the value of the primal is also 0 for any K . Therefore, given K , there is a feasible primal solution $\xi_i(a, s)$ that satisfies all the primal constraints, and CEB is satisfied. \square

Proposition 4.1. For $B \subset A$, the B -cone generated by unidentifiable deviation profiles is

$$\mathcal{K}(B) := \{ \lambda \geq 0 : \forall i \in I, a \in B, s \in S, \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = 0 \}.$$

By the Alternative Theorem (Rockafellar, 1970, p. 198), a given σ is implementable, i.e., there exists ζ to solve (***) , if and only if the following dual inequalities are satisfied:

$$\forall \lambda \in \mathcal{K}(\text{supp } \sigma), \quad \sum_{(i, a_i, \theta_i)} \lambda_i(a_i, \theta_i) \sum_{a_{-i}} \sigma(a) (v_i(b_i, a_{-i}) - v_i(a)) \leq 0.$$

In contrast, approximate implementability of σ as in Definition 3.1 is equivalent to the smaller system of inequalities indexed instead by $\lambda \in \mathcal{K}(A) \subset \mathcal{K}(\text{supp } \sigma)$. (Hence, exact implementability implies approximate). Now, if σ is completely mixed then $\sigma(a) > 0$ for all a , so $\mathcal{K}(\text{supp } \sigma) = \mathcal{K}(A)$. By DUD, $\mathcal{K}(A)$ consists of all $\lambda \geq 0$ with $\lambda_i(a_i, \theta_i) > 0$ implying $a_i = b_i$, where $\theta_i = (\rho_i, b_i)$. Therefore, $\sum_{(i, a_i, \theta_i)} \lambda_i(a_i, \theta_i) \sum_{a_{-i}} \sigma(a) (v_i(b_i, a_{-i}) - v_i(a)) = 0$, and implementability follows. For IOP, replacing $\mathcal{K}(B)$ with

$$\mathcal{K}_0(B) := \{ \lambda \geq 0 : \forall i \in I, a \in B, s \in S, \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i) (\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) = \eta(a, s) \}$$

leads to the corresponding result by an almost identical argument.

Clearly, the closure of the space of contracts satisfying (***) (and (**)) is contained in the space of contracts satisfying (*) (and (**)), so it remains only to show the converse containment. To this end, pick any (σ, ξ) satisfying (*) (and (**)). By the previous argument, the uniformly distributed correlated strategy with full support $\sigma^0 = (1/|A|, \dots, 1/|A|)$ is implementable (with budget balance). For any sequence of positive probabilities $\{p_m\}$ decreasing to 0, consider the sequence of contracts $\{(\sigma^m, \zeta^m)\}$ defined for every (i, a, s) by $\sigma^m(a) = p_m \sigma^0(a) + (1 - p_m) \sigma(a)$ and $\zeta_i^m(a, s) = p_m \zeta_i^0(a, s) + (1 - p_m) \xi_i(a, s) / \sigma^m(a)$. This sequence of contracts converges to (σ, ξ) and satisfies (***) (as well as (**)) for all m . \square

Proposition 4.4. We just prove the result with budget balance; the rest follows similarly. The dual of the metering problem of maximizing $\sum_a f(a)\sigma(a)$ subject to limited liability, approximate implementability, and budget balance is

$$\begin{aligned} V_f(v, \ell) &= \inf_{\lambda, \mu \geq 0, \eta, \kappa} \kappa \quad \text{s.t.} \\ \forall a \in A, \quad \kappa &\geq f(a) - \sum_{(i, \theta_i)} \lambda_i(a_i, \theta_i)(v_i(b_i, a_{-i}) - v_i(a)) - \sum_{(i, s)} \mu_i(a, s)\ell_i, \\ \forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) + \mu_i(a, s) &= \eta(a, s), \end{aligned}$$

where $\mu_i(a, s)$ is a multiplier on the liquidity constraint for player i at (a, s) . Adding the last family of equations with respect to s implies $\sum_s q_i(a, s) = \sum_s \eta(a, s)$ for every i . Therefore,

$$\sum_{(i, s)} \mu_i(a, s)\ell_i = \sum_{(i, s)} \eta(a, s)\ell_i = \sum_{s \in S} \eta(a, s)\widehat{\ell},$$

where $\widehat{\ell} = \sum_i \ell_i$, so we may eliminate $\mu_i(a, s)$ from the dual problem as follows:

$$\begin{aligned} V_f(v, \ell) &= \inf_{\lambda, \eta, \kappa} \kappa \quad \text{s.t.} \\ \forall a \in A, \quad \kappa &\geq f(a) - \sum_{(i, \theta_i)} \lambda_i(a_i, \theta_i)(v_i(b_i, a_{-i}) - v_i(a)) - \sum_{s \in S} \eta(a, s)\widehat{\ell} \\ \forall(i, a, s), \quad \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) &\leq \eta(a, s). \end{aligned}$$

Any two liability profiles ℓ and ℓ' with $\widehat{\ell} = \widehat{\ell}'$ lead to this same dual with the same value. \square

Proposition 4.5. The dual of the metering problem subject to participation is:

$$\begin{aligned} V_f(v) &= \inf_{\lambda, \pi \geq 0, \kappa, \eta} \kappa \quad \text{s.t.} \\ \forall a \in A, \quad \kappa &\geq f(a) - \sum_{(i, \theta_i)} \lambda_i(a_i, \theta_i)(v_i(b_i, a_{-i}) - v_i(a)) + \sum_{i \in I} \pi_i v_i(a) \\ \forall(i, a, s), \quad \pi_i \Pr(s|a) + \sum_{\theta_i \in \Theta_i} \lambda_i(a_i, \theta_i)(\Pr(s|a) - \Pr(s|\theta_i, a_{-i})) &= \eta(a, s) \end{aligned}$$

where π_i is a multiplier for player i 's participation constraint. Adding the second family of dual constraints with respect to $s \in S$, it follows that $\pi_i = \pi$ does not depend on i . Redefining $\eta(a, s)$ as $\eta(a, s) - \pi \Pr(s|a)$, the set of all feasible $\lambda \geq 0$ is the same as without participation constraints. Since $\sum_i v_i(a) \geq 0$ for all a , the dual is minimized by $\pi = 0$. \square

Lemma 5.3. If $a_i^m \rightarrow a_i$ then for any $g \in C(A_i)$,

$$\int_{A_i} g(b_i)[\beta_i(a_i^m)](db_i) = g(\beta(a_i^m)) \rightarrow g(\beta(a_i)),$$

since β is also continuous and composition preserves continuity. \square

Proposition 5.4. Sufficiency follows from Lemmata 5.2 and 5.3. For necessity, choose any $\lambda_i \in C(A_i, M(A_i))_+$. Without loss, $\lambda_i(A_i|a_i) = 1$. By continuity, a sequence $\{\lambda_i^m\}$ exists with λ_i^m simple (with finite range) for every m and $\lambda_i^m \rightarrow \lambda_i$, i.e., for every $g \in C(A_i)$, $\|\lambda_i^m - \lambda_i\|_g \rightarrow 0$. This implies that given m , $x_m := \int_A \int_{A_i} v_i(b_i, a_{-i}) \lambda_i^m(db_i|a_i) \sigma(da) = \sum_{k=1}^{K_i^m} \int_{A_{-i}} \int_{A_i} v_i(b_i, a_{-i}) \lambda_i^{m,k}(db_i) \sigma_i^{m,k}(da_{-i})$, where $\{\lambda_i^{m,k}\}_k$ is the range of λ_i^m and the measure $\sigma_i^{m,k}(\cdot) = \sigma(B_i^{m,k} \times \cdot)$ is defined with respect to the partition $\{B_i^{m,k} \subset A_i\}_k$ arising from the inverse of λ_i^m . By Fubini's Theorem, $x_m = \sum_{k=1}^{K_i^m} \int_{A_i} \int_{A_{-i}} v_i(b_i, a_{-i}) \sigma_i^{m,k}(da_{-i}) \lambda_i^{m,k}(db_i)$. Notice that the integrand $g_i^{m,k}(b_i) := \int_{A_{-i}} v_i(b_i, a_{-i}) \sigma_i^{m,k}(da_{-i})$ is continuous in b_i , so there exists $\beta_i^{m,k} \in A_i$ such that $g_i^{m,k}(\beta_i^{m,k}) = \int_{A_i} g_i^{m,k}(b_i) \lambda_i^{m,k}(db_i)$. This defines the deviation $\beta_i^m(a_i) = \beta_i^{m,k}$ if $a_i \in B_i^{m,k}$. By construction, $x_m = \int_A v_i(\beta_i^m(a_i), a_{-i}) \sigma(da)$, and since σ is a correlated equilibrium, this is less than or equal to $\int_A v_i(a) \sigma(da)$, so λ_i^m is unprofitable. Finally, by assumption $\int_A \int_{A_i} v_i(b_i, a_{-i}) \lambda_i^m(db_i|a_i) \sigma(da)$ converges to $\int_A \int_{A_i} v_i(b_i, a_{-i}) \lambda_i(db_i|a_i) \sigma(da)$ as $m \rightarrow \infty$, rendering λ_i also unprofitable. \square

Lemma 5.5. First we show that given $\sigma \in M(A)$, the functional $\lambda_i \mapsto F_i^*(\sigma)(\lambda_i)$ is continuous (linearity is obvious). Suppose $\lambda_i^m \rightarrow \lambda_i$ in $C(A_i, M(A_i))$, so for every $g \in C(A_i)$, $\|\lambda_i^m - \lambda_i\|_g \rightarrow 0$. (Appendix A defines this notation.) Since v_i is continuous, $|F_i^*(\sigma)(\lambda_i^m) - F_i^*(\sigma)(\lambda_i)| = |F_i^*(\sigma)(\lambda_i^m - \lambda_i)| \leq \|\lambda_i^m - \lambda_i\|_{f_i} \rightarrow 0$, where $f_i(a_i) = 2\|v_i\|$ is constant and finite, hence also continuous. It remains to prove that the operator F_i^* is continuous (linearity is again obvious). Suppose $\sigma_m \rightarrow \sigma$, i.e., for every $g \in C(A)$, $\int_A g(a) \sigma_m(da) \rightarrow \int_A g(a) \sigma(da)$. Again, since v_i is continuous, $|F_i^*(\sigma^m)(\lambda_i) - F_i^*(\sigma)(\lambda_i)| = |F_i^*(\sigma^m - \sigma)(\lambda_i)| \rightarrow 0$ because $g(a) = \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(db_i|a_i)$ is continuous. \square

Lemma 5.6. Since U is weakly increasing, it follows that, $U(f) - U(g) \leq U(f \vee g) - U(g)$ for any f and g in $C(A)$, where $f \vee g(a) := \max\{f(a), g(a)\}$ for every $a \in A$. Suppose that $\{(\lambda^m, \kappa^m) : m \in \mathbb{N}\}$ solves $U(g)$, i.e., (λ^m, κ^m) is g -feasible for every m and $\kappa^m \rightarrow U(g)$. For every $\varepsilon > 0$, pick m_ε such that $\kappa^{m_\varepsilon} - U(g) < \varepsilon$, and define $\widehat{\kappa}^\varepsilon$ as

$$\widehat{\kappa}^\varepsilon := \inf_{\kappa} \kappa \quad \text{s.t.} \quad \forall a \in A, \quad \kappa + \sum_{i=1}^n \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i^{m_\varepsilon}(db_i|a_i) \geq f \vee g(a).$$

Clearly, $U(f \vee g) \leq \widehat{\kappa}^\varepsilon$, so $U(f) - U(g) \leq \widehat{\kappa}^\varepsilon - \kappa^{m_\varepsilon} + \varepsilon$. To reduce notation, write $\Lambda_\varepsilon(a) = \sum_{i=1}^n \int_{A_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i^{m_\varepsilon}(db_i|a_i)$. For every a , both $\widehat{\kappa}^\varepsilon \geq f \vee g(a) - \Lambda_\varepsilon(a)$ and $\kappa^{m_\varepsilon} \geq g(a) - \Lambda_\varepsilon(a)$. By definition of $\widehat{\kappa}^\varepsilon$, given $\eta > 0$ there is \widehat{a} with $\widehat{\kappa}^\varepsilon - \eta < f \vee g(\widehat{a}) - \Lambda_\varepsilon(\widehat{a})$, so

$$\begin{aligned} \widehat{\kappa}^\varepsilon - \kappa^{m_\varepsilon} &\leq \widehat{\kappa}^\varepsilon - (g(\widehat{a}) - \Lambda_\varepsilon(\widehat{a})) < \eta + f \vee g(\widehat{a}) - \Lambda_\varepsilon(\widehat{a}) - (g(\widehat{a}) - \Lambda_\varepsilon(\widehat{a})) \\ &= \eta + f \vee g(\widehat{a}) - g(\widehat{a}) \leq \eta + \|f \vee g - g\| \leq \eta + \|f - g\|, \end{aligned}$$

where $\|\cdot\|$ stands for the sup norm. Hence, for any $\varepsilon > 0$ and $\eta > 0$, $U(f) - U(g) \leq \|f - g\| + \eta + \varepsilon$, so $U(f) - U(g) \leq \|f - g\|$. Reversing f and g proves the lemma. \square

Theorem 5.7. The first part of the theorem follows by definition of subdifferential and the fact that U is subdifferentiable. For the second, it is easy to see that V is convex. Since $V(0) = 0$, it follows that $\partial V(0) = \{\pi \in \mathbb{R}^n : \forall w \in \mathbb{R}^n, \pi \cdot w \leq V(w)\}$. Clearly, the set of correlated equilibrium payoffs is contained in the subdifferential, because feasible payoffs are attainable for any w . On the other hand, if a payoff vector $\pi \in \mathbb{R}^n$ belongs to the convex set $\partial V(0)$ then it may be attained by some correlated equilibrium. Indeed, for every $w \in \mathbb{R}^n$, let σ_w be a dual solution for the objective $\sum_i w_i v_i$. For every w , $\pi \cdot w \leq \int_A \sum_i w_i v_i(a) \sigma_w(da)$. By Caratheodory's Theorem (Rockafellar, 1970, p. 155), π is a convex combination of $n + 1$ vectors of the form $\int_A v(a) \sigma_{w_k}(da)$, where $v(a) = (v_1(a), \dots, v_n(a))$. Since the set of correlated equilibria is convex, π is attained as a correlated equilibrium payoff vector. \square

Theorem 5.11. We proceed in steps. Step 1 defines a family of linear programs indexed by right-hand side constraints. Step 2 shows that the value function of these problems has bounded steepness, so there is no duality gap between the primal and its dual. This justifies the use of duality. Step 3 then proves the theorem in line with the proof of Theorem 3.6.

– *Step 1: Defining the primal.* Consider the following linear program, called the *primal*.

$$\begin{aligned} V(g, h) &:= \inf_{\lambda \geq 0, \kappa} \kappa \quad \text{s.t.} \quad \lambda \in \prod_{i=1}^n C(A_i, M(\Theta_i)), \kappa \in \mathbb{R}, \\ \forall a \in A, \quad \kappa + \sum_{i=1}^n \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i|a_i) &\geq g(a), \text{ and} \\ \forall i \in I, a \in A, \varphi, \psi \in C_+(S), \\ - \int_S \varphi(s) h_i^-(ds, a) &\leq \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \lambda_i(d\theta_i|a_i), \\ \int_{\Theta_i} \int_S \psi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \lambda_i(d\theta_i|a_i) &\leq \int_S \psi(s) h_i^+(ds, a). \end{aligned}$$

In the notation of Appendix A, the spaces X and Y are given by

$$X = \mathbb{R} \times \prod_{i=1}^n C(A_i, M(\Theta_i)), \quad Y = C(A) \times \prod_{i=1}^{2n} C(A, M(S)).$$

– *Step 2: V has bounded steepness at $(f, 0)$.* In order to make meaningful use of the dual, we must show that no duality gap exists. By Theorem A.1, this amounts to showing that V has bounded steepness, i.e., given (g, h) and (fixed) f there is a constant $C \geq 0$ such that

$$\frac{V(f, 0) - V(g, h)}{\|(g, h) - (f, 0)\|} \leq C < +\infty.$$

We will attack this problem in parts by noticing first of all that

$$V(f, 0) - V(g, h) = \underbrace{V(f, 0) - V(g, 0)}_{U(f) - U(g)} + \underbrace{V(g, 0) - V(g, h)}_{W(0) - W(h)}.$$

– *Step 2a: The function U has bounded steepness, $|U(f) - U(g)| \leq \|f - g\|$.* This follows by exactly the same argument as for Lemma 5.6.

– *Step 2b: The function W has bounded steepness, too.* We will show that $W(0) - W(h) \leq \|h\|$. Clearly, W is a decreasing function, so without loss $h \geq 0$. For any $\varepsilon > 0$ and any pair (λ, κ) that is feasible for the primal with right-hand side constraints $(g, 0)$, there exists $(\lambda^\varepsilon, \kappa^\varepsilon)$ —feasible with respect to the constraints (g, h) —satisfying $W(h) > \kappa^\varepsilon - \varepsilon$. Hence, $W(0) - W(h) < \kappa - \kappa^\varepsilon + \varepsilon$. By feasibility, $\kappa \geq g(a) - \Lambda(a)$ for every $a \in A$, where $\Lambda(a) = \sum_i \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i|a_i)$. Similarly, $\kappa^\varepsilon \geq g(a) - \Lambda_\varepsilon(a)$ for every $a \in A$, where $\Lambda_\varepsilon(a)$ is the same as $\Lambda(a)$ but with λ^ε replacing λ . We will restrict attention to pairs (λ, κ) such that $\kappa = \sup_a \{g(a) - \Lambda(a)\}$. By continuity of g and Λ (inherited from the weak* continuity of λ , see Lemma 5.5) there exists $\hat{a} \in A$ such that $\kappa = g(\hat{a}) - \Lambda(\hat{a})$. For this \hat{a} , it follows that $W(0) - W(h) \leq \varepsilon + g(\hat{a}) - \Lambda(\hat{a}) - g(\hat{a}) + \Lambda_\varepsilon(\hat{a}) = \varepsilon - \Lambda(\hat{a}) + \Lambda_\varepsilon(\hat{a})$.

We will now bound $\Lambda_\varepsilon(\hat{a}) - \Lambda(\hat{a})$. Notice that for every i and a_i such that λ^ε satisfies

$$\sup_{\varphi \in C(S)_+} \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \lambda_i^\varepsilon(d\theta_i|a_i) = 0,$$

we may choose $\lambda_i(a_i) = \lambda_i^\varepsilon(a_i)$, so $\Lambda_\varepsilon(\hat{a}) - \Lambda(\hat{a}) = \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) (\lambda_i^\varepsilon - \lambda_i)(d\theta_i|a_i)$ can be made equal to zero, leaving nothing to prove. So, suppose the supremum is not zero.

For every $\varepsilon > 0$, rewrite $\lambda_i^\varepsilon = \alpha_i^\varepsilon \mu_i^\varepsilon$, where $\alpha_i^\varepsilon : A_i \rightarrow \Delta(\Theta_i)$ is a continuous deviation plan and $\mu_i^\varepsilon : A_i \rightarrow [0, \infty)$ is some scaling factor. Similarly, rewrite $\lambda_i = \alpha_i \mu_i$. Without loss of generality, $\mu_i = \mu_i^\varepsilon$. Indeed, if $\mu_i(a_i) > \mu_i^\varepsilon(a_i)$ for some $a_i \in A_i$ then let

$$\hat{\alpha}_i^\varepsilon(a_i) = \frac{\mu_i^\varepsilon(a_i)}{\mu_i(a_i)} \alpha_i^\varepsilon(a_i) + \left(1 - \frac{\mu_i^\varepsilon(a_i)}{\mu_i(a_i)}\right) [a_i],$$

where $[a_i]$ is Dirac measure, and let $\hat{\alpha}_i(a_i) = \alpha_i(a_i)$. If $\mu_i(a_i) < \mu_i^\varepsilon(a_i)$ for some a_i let

$$\hat{\alpha}_i(a_i) = \frac{\mu_i(a_i)}{\mu_i^\varepsilon(a_i)} \alpha_i(a_i) + \left(1 - \frac{\mu_i(a_i)}{\mu_i^\varepsilon(a_i)}\right) [a_i]$$

and $\hat{\alpha}_i^\varepsilon(a_i) = \alpha_i^\varepsilon(a_i)$. Let $\hat{\mu}_i(a_i) = \max\{\mu_i(a_i), \mu_i^\varepsilon(a_i)\}$, $\hat{\lambda}_i = \hat{\alpha}_i \hat{\mu}_i$, and $\hat{\lambda}_i^\varepsilon = \hat{\alpha}_i^\varepsilon \hat{\mu}_i^\varepsilon$. By construction, $\int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i|a_i) = \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \hat{\lambda}_i(d\theta_i|a_i)$ and $\int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \lambda_i(d\theta_i|a_i) = \int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \hat{\lambda}_i(d\theta_i|a_i)$; similarly for λ_i^ε and $\hat{\lambda}_i^\varepsilon$. Hence,

$$\Lambda_\varepsilon(a) - \Lambda(a) = \sum_{i=1}^n \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) (\hat{\alpha}_i^\varepsilon - \hat{\alpha}_i)(d\theta_i|a_i) \hat{\mu}_i(a_i).$$

By feasibility, if $\int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \hat{\alpha}_i^\varepsilon(d\theta_i|a_i) \neq 0$ then for every $\varphi \in C(S)$ such that the denominator below is not zero, and every a_{-i} ,

$$\hat{\mu}_i(a_i) \leq \frac{\min\{|\int_S \varphi(s) h_i^+(ds, a)|, |\int_S \varphi(s) h_i^-(ds, a)|\}}{\left| \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \hat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right|}$$

$$\begin{aligned}
&\leq \inf_{\varphi, a_{-i}} \frac{\min\{|\int_S \varphi(s) h_i^+(ds, a)|, |\int_S \varphi(s) h_i^-(ds, a)|\}}{\left| \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right|} \\
&\leq \inf_{\varphi, a_{-i}} \frac{\max\{|\int_S \varphi(s) h_i^+(ds, a)|, |\int_S \varphi(s) h_i^-(ds, a)|\}}{\left| \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right|} \\
&\leq \frac{\sup_{\varphi, a_{-i}} \{ \max\{|\int_S \varphi(s) h_i^+(ds, a)|, |\int_S \varphi(s) h_i^-(ds, a)|\} : \|\varphi\| = \sup_s |\varphi(s)| = 1 \}}{\sup_{\varphi, a_{-i}} \left\{ \left| \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right| : \|\varphi\| = 1 \right\}} \\
&\leq \frac{\sup_{\varphi, a} \{ \max\{|\int_S \varphi(s) h_i^+(ds, a)|, |\int_S \varphi(s) h_i^-(ds, a)|\} : \|\varphi\| = \sup_s |\varphi(s)| = 1 \}}{\sup_{\varphi, a_{-i}} \left\{ \left| \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right| : \|\varphi\| = 1 \right\}} \\
&= \frac{\max\{\|h_i^+\|, \|h_i^-\|\}}{\sup_{a_{-i}} \left\| \int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right\|} \leq \|h_i\| / c_i,
\end{aligned}$$

where $\|h_i\|$ is defined as $\max\{\|h_i^+\|, \|h_i^-\|\}$; the last inequality follows by Definition 5.9.

For every $\varepsilon > 0$ and $a \in A$, if $\sup_{a_{-i}} \left\| \int_{\Theta_i} \Pr(a) - \Pr(\theta_i, a_{-i}) \widehat{\alpha}_i^\varepsilon(d\theta_i|a_i) \right\| \neq 0$ then

$$\begin{aligned}
\Lambda_\varepsilon(a) - \Lambda(a) &\leq \sum_{i=1}^n \left| \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) (\widehat{\alpha}_i^\varepsilon - \widehat{\alpha}_i)(d\theta_i|a_i) \right| \widehat{\mu}_i(a_i) \\
&\leq \sum_{i=1}^n \left| \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) (\widehat{\alpha}_i^\varepsilon - \widehat{\alpha}_i)(d\theta_i|a_i) \right| \|h_i\| / c_i \leq \sum_{i=1}^n \frac{2\|v_i\|}{c_i} \|h_i\|.
\end{aligned}$$

Now, h lives in the space $\prod_i C(S, M(A))^*$, and $\|h\| = \sum_i \frac{2\|v_i\|}{c_i} \|h_i\|$ clearly defines a norm on that space. Therefore, W is bounded by a norm of h , as claimed.

– *Step 2c: Steps 2a and 2b imply that V has bounded steepness.* We may define the norm of any (g, h) in the domain of V by $\|(g, h)\| = \|g\| + \|h\|$, whence

$$V(g, h) - V(f, 0) \leq \|g - f\| + \|h\| = \|(g - f, h)\| = \|(g, h) - (f, 0)\|.$$

By Theorem A.1, the *dual* (below) when $h = 0$ has a solution and there is no duality gap.

$$\begin{aligned}
&\sup_{\sigma \geq 0, \xi} \int_A f(a) \sigma(da) \quad \text{s.t.} \quad \forall i \in I, \lambda_i \in C(A_i, M(\Theta_i))_+, \\
&\int_A \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i|a_i) \sigma(da) \leq \\
&\int_{A \times S} \int_{\Theta_i} \Pr(ds|a) - \Pr(ds|\theta_i, a_{-i}) \lambda_i(d\theta_i|a_i) \xi_i(s, da).
\end{aligned}$$

When $h = 0$, the dual incentive constraints coincide with (\dagger) .

– *Step 3: Proof of the theorem via duality.* For necessity, if DUD holds then the only feasible λ in the primal with left-hand side constraints $(f, 0)$ satisfy $\text{supp } \lambda_i(a_i) \subset \{a_i\} \times R_i$,

so $V(f, 0) = \max\{f(a) : a \in A\}$ for every f in the domain of V . In particular, f can be picked to make any pure action profile approximately implementable with some contract. For sufficiency, if DUD fails then there is an undetectable deviation profile involving disobedience at some action profile a . As in the finite case, if a_i is strictly dominated for some player i by this disobedience then a is not approximately implementable. \square

Theorem 5.13. We follow the steps of Theorem 5.11 with slightly different linear programs.

– *Step 1: Defining the primal.* Consider the following problem called the *primal*.

$$\begin{aligned}
V(g, h) &:= \inf_{\lambda \geq 0, \eta, \kappa} \kappa \quad \text{s.t.} \quad \lambda \in \prod_{i=1}^n C(A_i, M(\Theta_i)), \quad \eta \in C(A, M(S)), \quad \kappa \in \mathbb{R}, \\
&\forall a \in A, \quad \kappa + \sum_{i=1}^n \int_{\Theta_i} v_i(b_i, a_{-i}) - v_i(a) \lambda_i(d\theta_i|a_i) \geq g(a), \quad \text{and} \\
&\forall i \in I, a \in A, \varphi, \psi \in C_+(S), \\
&\int_S \varphi(s) (\eta(ds, a) - h_i^-(ds, a)) \leq \int_{\Theta_i} \int_S \varphi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \lambda_i(d\theta_i|a_i), \\
&\int_{\Theta_i} \int_S \psi(s) (\Pr(ds|a) - \Pr(ds|\theta_i, a_{-i})) \lambda_i(d\theta_i|a_i) \leq \int_S \psi(s) (\eta(ds, a) + h_i^+(ds, a)).
\end{aligned}$$

In the notation of the previous subsection, the spaces X and Y are given by

$$X = \mathbb{R} \times C(A, M(S)) \times \prod_{i=1}^n C(A_i, M(\Theta_i)), \quad Y = C(A) \times \prod_{i=1}^{2n} C(A, M(S))^*.$$

– *Step 2: V has bounded steepness at $(f, 0)$.* To justify the use of duality in this more complicated context we must show that V as defined here has bounded steepness. Just as for Theorem 5.11, let us break up $V(g, h) - V(f, 0)$ into two functions, U and W . By the same argument as for Lemma 5.6, U has bounded steepness, so it only remains to bound the steepness of W . This is just like Step 2b in the proof of Theorem 5.11 except for the following alteration. Instead of choosing any (λ, η, κ) that is feasible for the primal with right-hand side constraints $(f, 0)$, for every $\varepsilon > 0$, choose $(\lambda^{(\varepsilon)}, \kappa^{(\varepsilon)})$ to satisfy the primal constraints with η^ε chosen to be feasible for the problem with right-hand side constraints $(\lambda^\varepsilon, \eta^\varepsilon, \kappa^\varepsilon)$ and $\kappa^\varepsilon < W(h) + \varepsilon$. Step 2 now follows.

– *Step 3: Proof of the theorem via duality.* Same argument as Theorem 5.11. \square

Corollary 5.16. First notice that the repeated game is continuous. Define the action space of player i to be Θ^i , and let the mediator be player 0 with action space \mathcal{A} . Define the following topology on a player's action space: an open set is any set of actions which are identical at all possible partial histories of length at most t for $t = 1, 2, \dots$. Every player's

action space is easily seen to be compact, Hausdorff. With the product topology on the space of action profiles, each player's discounted average payoff is clearly continuous (the mediator's payoff is identically zero). Hence, we may apply Theorem 5.7 to the dual of the family of linear programs below (a "metamediator" is not necessary), indexed by $w \in \mathbb{R}^n$:

$$\begin{aligned} & \sup_{\sigma \geq 0} (1 - \delta) \sum_{\mathbf{a} \in \mathcal{A}} \sigma(\mathbf{a}) \sum_{(i,t,h_t)} \delta^{t-1} w_i v_i(\mathbf{a}(h_t)) \Pr(h_t|\mathbf{a}) \text{ s.t. } \sum_{\mathbf{a} \in \mathcal{A}} \sigma(\mathbf{a}) = 1, \text{ and } \forall(i, \theta^i), \\ & \sum_{\mathbf{a} \in \mathcal{A}} \sigma(\mathbf{a}) \sum_{(t,h_t,\hat{h}_t^i)} \delta^{t-1} [v_i(b^i(\hat{h}_t^i, \mathbf{a}_i(h_t)), \mathbf{a}_{-i}(h_t)) \Pr(h_t, \hat{h}_t^i|\theta^i, \mathbf{a}) - v_i(\mathbf{a}(h_t)) \Pr(h_t|\mathbf{a})] \leq 0. \end{aligned}$$

Since this is the dual of the problem defining $V_\delta(w)$, the result now follows. \square

Theorem 5.18. We proceed in four steps. The first step is to describe a more tractable version of the theorem. In the second step we derive a result that will be useful in the third step, where we prove the version of the theorem presented in the first step when $w_i > 0$ for some i . Finally, the fourth step proves the folk theorem including when $w \leq 0$.

– *Step 1: Another description of the folk theorem.* Given w , let $a(w) \in A$ be an action profile that maximizes $\sum_i w_i v_i(a)$. Step 3 below will show that $\lim_{T \rightarrow \infty} V_1^T(w) = \sum_i w_i v_i(a)$ for all $w \geq 0$ by solving the dual problem below whose value clearly equals $V_1^T(w)$:

$$\begin{aligned} V_1^T(w) &= \inf_{\lambda, \eta \geq 0, \kappa} \kappa \text{ s.t. } \forall \mathbf{a}_T \in \mathcal{A}_T, \kappa \geq \frac{1}{T} \left\{ \sum_{(i,t,h_t)} w_i v_i(\mathbf{a}_T(h_t)) \Pr(h_t|\mathbf{a}_T) - \right. \\ & \left. \sum_{(i,\theta^i)} \lambda_i(\theta^i) \sum_{(t,h_t,\hat{h}_t^i)} v_i(b^i(\hat{h}_t^i, \mathbf{a}_i(h_t)), \mathbf{a}_{-i}(h_t)) \Pr(h_t, \hat{h}_t^i|\theta^i, \mathbf{a}) - v_i(\mathbf{a}(h_t)) \Pr(h_t|\mathbf{a}) \right\}, \\ \forall(i, \mathbf{a}_T, h_{T+1}), \sum_{\theta_T^i} \lambda_i(\theta_T^i) (\Pr(h_{T+1}|\theta_T^i, \mathbf{a}_T) - \Pr(h_{T+1}|\mathbf{a}_T)) &= w_i (\Pr(h_{T+1}|\mathbf{a}_T) - \eta(h_{T+1})), \end{aligned}$$

where $\eta : H_{T+1} \rightarrow \mathbb{R}_+$ is a multiplier for the constraint $\sum_i w_i \xi_i(h_{T+1}) \leq 0$.

Let $\mathbf{a}_T(w)$ be the communication strategy that assigns $a(w)$ after every history. Clearly, $V_1^T(w) \leq \sum_i w_i v_i(a(w))$, since $\lambda = 0$ and $\eta(h_{T+1}) = \Pr(h_{T+1}|\mathbf{a}_T(w))$ are feasible. We want to show that $V_1^T(w) \geq \sum_i w_i v_i(a(w)) - o(T)$, where $o(T)$ tends to 0 as T explodes.

– *Step 2: If \Pr satisfies SCI then for every $i \in I$, $\alpha_{-i} \in \Delta(A_{-i})$ completely mixed, $\gamma > 0$ and $\gamma' > \gamma$, there exists $\phi_i : A \times S \rightarrow \mathbb{R}$ such that*

$$\begin{aligned} \forall(a_i, s_i), \sum_{(a_{-i}, s_{-i})} \phi_i(a, s) \Pr(a_{-i}, s_{-i}|a_i, \alpha_{-i}, s_i) &= \gamma \\ \forall(b_i, t_i) \neq (a_i, s_i), \sum_{(a_{-i}, s_{-i})} \phi_i(a, s) \Pr(a_{-i}, s_{-i}|b_i, \alpha_{-i}, t_i) &\geq \gamma'. \end{aligned}$$

This follows immediately from the Alternative Theorem (Rockafellar, 1970, p. 198).

Let W_ϵ be the set of w such that $\|w\| = 1$ and $w_i \geq \epsilon$ for some $i \in I$.

– Step 3: If \Pr satisfies SCI then given $z > 0$ and $\epsilon > 0$, there exists T' such that for all $T \geq T'$ and $w \in W_\epsilon$, $V_1^T(w) \geq \sum_i w_i v_i(a(w)) - z$. Let $V^*(w) = \sum_{i=1}^n w_i v_i(a(w))$. Given $\mathbf{a}_T(w)$, the first term of the inequality constraint (in the problem of Step 1) is $V^*(w)$. We will show that the second term becomes negligible, uniformly with respect to $w \in W_\epsilon$ as $T \rightarrow \infty$, for any feasible λ and η . This proves that V_1^T cannot be much smaller than $V^*(w)$, and therefore also proves the step because $V^*(w)$ itself is feasible by using $\mathbf{a}_T(w)$ and setting $\lambda = 0$, $\eta(h_{T+1}) = \Pr(h_{T+1}|\mathbf{a}_T(w))$.

Let i be any individual with $w_i \in [0, 1]$. From the equality constraint in Step 1, we have

$$\forall \mathbf{a}_T \in \mathcal{A}_T, \quad \sum_{\theta_T^i} \lambda_i(\theta_T^i) \sum_{\hat{h}_{T+1}^i} (\Pr(h_{T+1}, \hat{h}_{T+1}^i | \theta_T^i, \mathbf{a}_T) - \Pr(h_{T+1} | \mathbf{a}_T)) \leq \Pr(h_{T+1} | \mathbf{a}_T). \quad (T)$$

With a slight abuse of notation, let α_T be a T -period communication strategy recommending $a^i(w)$ to player i and a completely mixed correlated strategy $\alpha^{-i} \in A_{-i}$ to the others after every history. (T) holds for all \mathbf{a}_T , hence also for α_T . By definition of \Pr ,

$$\Pr(h_{T+1}, \hat{h}_{T+1}^i | \theta_T^i, \alpha_T) = \Pr(h_T, \hat{h}_T^i | \theta_T^i, \alpha_T) \Pr(\hat{s}_T^i | \hat{a}_T^i, \alpha^{-i}) \Pr(a_T^{-i}, s_T^{-i} | \hat{a}_T^i, \alpha^{-i}, \hat{s}_T^i)$$

if $a_T^i = a^i(w)$, $s_T^i = \rho_T^i(\hat{h}_T^i)$ and $\hat{a}_T^i = b^i(\hat{h}_T^i, a^i(w))$; otherwise $\Pr(h_{T+1}, \hat{h}_{T+1}^i | \alpha_T, \theta_T^i) = 0$.

Now multiply both sides of (T) by $\phi_i(a^i(w), a_T^{-i}, s_T)$ from Step 2 and sum them with respect to both (a_T^{-i}, s_T^{-i}) and \hat{s}_T^i . This yields

$$\begin{aligned} \sum_{\theta_T^i} \lambda_i(\theta_T^i) \sum_{\hat{h}_T^i} \Pr(h_T, \hat{h}_T^i | \theta_T^i, \alpha_T) \sum_{\hat{s}_T^i \in S_i} \Pr(\hat{s}_T^i | \hat{a}_T^i, \alpha^{-i}) \gamma_i(\theta_T^i, k_T^i, \hat{s}_T^i) - \Pr(h_T | \alpha_T) \gamma \\ \leq \Pr(h_T | \alpha_T) \gamma. \end{aligned}$$

The value of $\gamma_i(\theta_T^i, k_T^i, t_T^i)$ is γ when $b^i(k_T^i) = a^i(w)$ and player i tells the truth, i.e. $\rho^i(k_T^i, b^i(k_T^i), t_T^i) = t_T^i$. It is at least as large as γ' otherwise. Repeating this algorithm for $\phi_i(a^i(w), a_t^{-i}, s_t)$ and $t = T-1, T-2, \dots$, and dividing both sides by γ^T , we obtain:

$$\begin{aligned} 1 &\geq \sum_{\theta_T^i} \lambda_i(\theta_T^i) \left[\sum_{h_{T+1}^i} \Pr(h_{T+1}^i | \alpha_T, \theta_T^i) \left(\frac{\gamma'}{\gamma} \right)^{l_i(\theta_T^i, h_{T+1}^i)} - 1 \right] \\ &\geq \sum_{\theta_T^i} \lambda_i(\theta_T^i) \left[\left(\frac{\gamma'}{\gamma} \right)^{El_i(\theta_T^i)} - 1 \right] \end{aligned}$$

where $l_i(\theta_T^i, h_{T+1}^i)$ represents the number of periods in which player i either deviated from $a^i(w)$ and/or lied along h_{T+1}^i by playing θ_T^i and $El_i(\theta_T^i)$ is its expected value. Let $\Theta^i(k)$ be the set of player i 's strategy for which $El_i(\theta_T^i)$ is larger than or equal to k . Then the above inequality implies that

$$\sum_{\theta_T^i \in \Theta^i(k)} \lambda_i(\theta_T^i) \leq \frac{1}{\left(\frac{\gamma'}{\gamma} \right)^k - 1}$$

Let $\bar{v} = \max_{(i,a,a')} |v_i(a) - v_i(a')|$, to be used as an upper bound for deviation gains. Given $\mathbf{a}_T(w)$, the second term of the dual inequalities for player i is now bounded as follows:

$$\begin{aligned} \frac{1}{T} \sum_{\theta_T^i} \lambda_i(\theta_T^i) \left\{ \sum_{(t,h_t,\hat{h}_t^i)} v_i(b^i(\hat{h}_t^i, a^i(w)), a^{-i}(w)) \Pr(h_t, \hat{h}_t^i | \theta_T^i, \mathbf{a}_T(w)) - v_i(a^i(w)) \Pr(h_t | \mathbf{a}_T(w)) \right\} \\ \leq \frac{1}{T} \left(\frac{1}{\frac{\gamma'}{\gamma} - 1} k \bar{v} + \frac{1}{\left(\frac{\gamma'}{\gamma}\right)^k - 1} \bar{v} \right). \end{aligned}$$

The second line follows because $\sum_{\theta_T^i} \lambda_i(\theta_T^i) \leq 1/(\frac{\gamma'}{\gamma} - 1)$ for strategies that deviate at most $k - 1$ times in expectation and $\sum_{\theta_T^i} \lambda_i(\theta_T^i) \leq 1((\frac{\gamma'}{\gamma})^k - 1)$ for strategies that deviate k times or more in expectation. This can be made arbitrarily close to 0 as taking k large and letting $T \rightarrow \infty$, regardless of $w^i \in [0, 1]$ and $i \in I$.

Next, consider an individual j with negative weight $w_j < 0$. Without loss of generality, suppose $w_1 \geq \epsilon$, so the dual equality constraint implies

$$\begin{aligned} \sum_{\theta_T^j} \lambda_j(\theta_T^j) (\Pr(h_{T+1} | \theta_T^j, \mathbf{a}_T) - \Pr(h_{T+1} | \mathbf{a}_T)) \\ = \frac{w_j}{w_1} \sum_{\theta_T^1} \lambda_1(\theta_T^1) (\Pr(h_{T+1} | \theta_T^j, \mathbf{a}_T) - \Pr(h_{T+1} | \mathbf{a}_T)) \leq \frac{1}{\epsilon} \sum_{\theta_T^1} \lambda_1(\theta_T^1) \Pr(h_{T+1} | \mathbf{a}_T). \end{aligned}$$

Using $\phi_j(a^j(w), a_T^{-j}, s_T)$ from Step 2 as before, we find

$$\sum_{\theta_T^j \in \Theta^j(k)} \lambda_j(\theta_T^j) \left[\left(\frac{\gamma'}{\gamma}\right)^{El_j(\theta_T^j)} - 1 \right] \leq \frac{1}{\epsilon} \sum_{\theta_T^1} \lambda_1(\theta_T^1) \leq \frac{1}{\epsilon} \left(\frac{1}{\left(\frac{\gamma'}{\gamma}\right)^k - 1} + \frac{1}{\frac{\gamma'}{\gamma} - 1} \right).$$

Therefore we obtain a bound for $\sum_{\theta_T^j \in \Theta^j(k)} \lambda_j(\theta_T^j)$:

$$\sum_{\theta_T^j \in \Theta^j(k)} \lambda_j(\theta_T^j) \leq \frac{1}{\left(\frac{\gamma'}{\gamma}\right)^k - 1} \frac{1}{\epsilon} \left(\frac{1}{\left(\frac{\gamma'}{\gamma}\right)^k - 1} + \frac{1}{\frac{\gamma'}{\gamma} - 1} \right)$$

which also converges to 0 as $k \rightarrow \infty$. It follows that the second term of the inequality constraint for player j is also bounded by an arbitrarily small positive number after taking k large and letting $T \rightarrow \infty$. This proves the step.

– *Step 4: Proving the Folk Theorem.* Now we can prove a folk theorem with respect to payoffs from any stage-game correlated equilibrium, σ^* . By Step 3, $V^*(w)$ can be approximated by $V_1^T(w)$ uniformly across $w \in W_\epsilon$ as $T \rightarrow \infty$. Since σ^* is always sustainable without transfers, $V_1^T(w) \geq \sum_{(i,a)} \sigma^*(a) w_i v_i(a)$ always holds for negative w . For any smooth compact convex subset in the interior of $V(\sigma^*)$, take z and ϵ small enough and T large enough so that the set is contained in $\text{int } \partial V_1^T(0)$. By Lemma 5.17, it is also included in $\liminf_{\delta \rightarrow 1} E_\delta^T$. Therefore, $V(\sigma^*) \subset \liminf_{T \rightarrow \infty} \liminf_{\delta \rightarrow 1} E_\delta^T$. \square