Optimum Contracts with Public and Private Monitoring

David Rahman, UCSD*

November 6, 2005

Abstract

This paper develops a formal approach to characterize optimum contracts in general economic environments. Convex independence of a public as well as private monitoring technology is identified as both necessary and sufficient for any team to achieve its first best outcome, making essential use of the duality theory of linear programming. With public monitoring, optimum contracts are a version of Holmström's (1982) team punishments, and team outcomes are implemented in Aumann's (1974) correlated equilibrium. With private monitoring, so-called loyalty-testing contracts are generally required to provide adequate incentives for monitors to both make the appropriate monitoring effort and truthfully report monitoring signals, and team outcomes are implemented in a version of Myerson's (1986) sequential communication equilibrium.

JEL Classification: D21, D23, D82.
Keywords: contracts, private monitoring, incentives, correlated equilibrium, sequential communication equilibrium, convex independence, duality.

*Please send any comments to drahman@euclid.ucsd.edu. This paper is based on Chapter 4 of my doctoral dissertation. I thank Joe Ostroy for invaluable guidance and Harold Demsetz for delightful conversations during the initial stages of this research. I am also grateful to David Levine and Bill Zame for encouragement and insightful advice, as well as Christian Hellwig, Ichiro Obara, Vasiliki Skreta, and participants at the UCLA theory proseminar for useful discussions.

1 Introduction

The economic organization of a team is divided into three components: individual actions, the team's trades, and the allocation of information to team members.

Individual actions may be further categorized into three kinds: strictly productive activities, monitoring behavior, and reporting. Productive actions could be thought of as a form of working, such as digging a mine. Monitoring may be thought of as collecting evidence regarding the amount of work incurred by workers, or simply observing them. Reporting may or may not be verifiable. Workers are compensated for their effort with incentive schemes that reward them contingent on their monitor's report. Monitors who also report, on the other hand, must be rewarded with "loyalty-testing" contracts in order that they make the required monitoring effort and report truthfully their monitoring signals. Such loyalty-testing contracts take the following form. A disinterested, third party correlates his recommendation of effort to workers with report-contingent contracts for the monitor. Thus, the monitor does not know the contract he faces unless he incurs the effort of monitoring workers. Truthful reporting is then guaranteed by asking the monitor to confirm the third party's recommendations to workers.

A team's trades play two important roles. Firstly, they add value by improving the team's welfare. Secondly, they may relax incentive constraints, in other words, they may make certain individual actions more desirable. Contractual payments are included in the trades of a team.

As for allocating information, there are three aspects relevant to the team. Firstly, there is information regarding the behavior of others. For instance, the team might be better off if some individuals are left uncertain of the actions of other team members. Secondly, there is information regarding the team's trade. For the same reasons as with behavioral information, it may be best for the team if some or all individuals remain unaware of the team's trading strategy. Finally, there may be payoff-relevant private information acquired prior to the formation of a team. Some of this information may be revealed to prospective team members and some of it may remain withheld after the team has formed. This last form of private information, although admittedly important for understanding organizations, is not studied here.[1]

[1] It is studied in Rahman (2005a), where it is argued that optimum contracts induce information revelation before team formation via *residual claims*, and after team formation via *control rights*.

This paper studies how teams can achieve economic organization after having formed. A team's choice of organizational design is determined to a large extent by its ability to provide appropriate incentives to team members with contractual rewards. Two important examples are the allocation of monitoring responsibility and information about the team's behavior. According to Alchian and Demsetz (1972, page 778, their footnote):

> *Two key demands are placed on an economic organization—metering input productivity and metering rewards.*[2]

Such view of economic organization is formalized below. The problem of measuring input productivity is interpreted as one of monitoring, and the problem of apportioning it as one of designing adequate incentive contracts. The problem of apportioning rewards is understood as arising from the unverifiability of monitoring, and the problem of measuring rewards as one of creating adequate incentives for certain team members to exert monitoring effort.

The paper is divided into two parts. The first part examines the problem of metering input productivity and the second part is concerned with metering rewards.

In Section 3 we describe a team that is subject to *public monitoring*, meaning that the team will be able to condition trades not just on effort recommendations to players but also on the realization of a public signal whose distribution will generally depend on team members' behavior. Thus, the problem of measuring input productivity is mitigated by having some members exert monitoring effort. The problem of controlling such productivity is solved by the team's contractual strategy which will typically involve incentive rewards and punishments.

In Section 4 we study a team that is subject to *private monitoring*, meaning that signal realizations are privately observed by team members who will then have the opportunity to report them dishonestly to the rest of the team. The problem of apportioning rewards is solved by inducing monitors to reveal their private information truthfully, accomplished by rewarding them with recommendation-contingent contractual payments. The problem of measuring rewards is solved by the initial uncertainty imposed on monitors' incentive scheme so that they prefer exerting an optimal amount of monitoring effort.

[2] *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.*

In conclusion, team punishments characterize contractual payments with public monitoring, and loyalty-testing contracts characterize payments with private monitoring.

The use of duality to analyze optimum contracts is prevalent throughout the paper (especially Sections 3.4 and 4.4). We extend the linear programming approach of Nau and McCardle (1990) and Myerson (1997), who used duality to show existence as well as consider refinements of Aumann's (1974) correlated equilibrium, to the current contract-theoretical setting with useful economic insights.

Specifically, we find the weakest sufficient condition, called *convex independence*, for implementability of any team's first-best outcomes (Theorems 3.10, 4.13, and 4.20).[3] Convex independence of a public or private monitoring technology means that every individual action can be statistically distinguished from any possible deviation (be it pure or mixed). Interpreting dishonesty as a form of disobedience, optimum contracts with private monitoring are a natural extension of their public counterpart.

The motivating examples of Section 2 with Robinson (a reporting monitor) and Friday (a worker) answer many organizational questions, like Holmström's (1982, page 339):

> ... *what determines the choice of monitors; and how should output be shared so as to provide all members of the organization* (*including monitors*) *with the best incentives to perform?*

As it turns out, it is important to not only consider how output is allocated, but also how information is shared to provide members of the organization with incentives to perform. As Robinson and Friday show, it is necessary that Friday's behavior creates uncertainty on Robinson (as part of the aforementioned loyalty-testing contracts) in order for the team to face the right incentives. Also, to save on monitoring effort, it is useful that Friday be kept in the dark regarding Robinson's behavior.

Therefore, how output is shared to provide incentives largely determines the choice of monitors. Monitoring is intrinsically unproductive, its only value comes from expanding contractual opportunities to provide incentives to workers. With limited scope for rewards (see Example 4.8), individual preferences may also determine the choice of monitors. For instance, a monitor who prefers to report deviations when they occur might lower the team's opportunity cost of monitoring. (See Section 4.3.)

[3]This condition is related to the literature on repeated games in Section 5, where the paper concludes with comments and connections.

Another question answered by this model is Alchian and Demsetz's (1972, page 782):

> *But who will monitor the monitor?*

According to them, the monitor monitors himself if he's made residual claimant.[4] Contrariwise, we propose that it is the principal who monitors the monitor by "testing his loyalty." He secretly recommends Friday to occasionally shirk, so that Robinson prefers to monitor when facing recommendation-contingent rewards. In fact, the team is better off if Robinson is not the principal so that he may face such uncertainty regarding Friday's effort. Robinson must not know what the principal knows (his recommendation to Friday).

A similar argument was put forward by Strausz (1997), who argues that delegated monitoring dominates monitoring by a principal who—when facing a budget-balance constraint—cannot commit to the agent that he will verify the agent's effort when it is only privately verifiable. However, Strausz (1997) assumes that monitoring signals are "hard evidence," so that Robinson cannot misreport his private signal, rendering loyalty tests unnecessary. With "soft evidence," loyalty tests become necessary.

These arguments confirm Holmström's (1982, page 325) view on the main economic role of the principal:

> *... the principal's role is not essentially one of monitoring ... the principal's primary role is to break the budget-balance constraint.*

Our model agrees with and adds to Holmström's view. Firstly, team members are given incentives to perform via standard incentive contracts (e.g., team punishments) or loyalty-testing contracts, without having to give any one team member claims to the team's residual. Secondly, the mediating principal (or as we call him, the organizer) is not essentially a monitor. Moreover, the team is better off if he isn't, as was already argued. By construction, our model does not show the breaking of budget-balance explicitly (since we only impose *expected* budget balance), but the same effect that Holmström alludes to is present here in that signal-contingent payments to individuals may differ across signal realizations. Finally, our model adds to Holmström's comment by arguing that, in a private monitoring environment, the principal's role includes mediating workers as a crucial part of incentive provision.

[4] We think of a residual claimant as a team member whose payoff depends on the team's profit.

2 Robinson and Friday

In this section we study two related examples involving contractual variations on a three-player game that attempts to typify the contractual relationship between a principal, an agent, and a monitor. The first player is Robinson, who can either monitor or shirk. The second player is Friday, who can either work or shirk. The third player is a so-called mediating principal, a disinterested third party who makes recommendations to the other players and enforces contingent contractual payments.

Suppose that the principal's utility is constant regardless of the outcome of the game, and consider the following two-player game played by Robinson (the row player) and Friday (the column player), who interact according to the bi-matrix below.

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

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.

It may be possible, by way of contract, to convince Friday to work, as long as there is a way of rewarding him when he works and/or punishing him when he shirks. To this end, we assume that the team's contractual technology is given by a set $S = \{g, b\}$ (so there are only two possible signal realizations contingent upon which contracts may be written) together with the conditional probability system given by

$$\begin{aligned}
\Pr(g|\text{monitor,work}) &= 1 = 1 - \Pr(b|\text{monitor,work}) \\
\Pr(b|\text{monitor,shirk}) &= 1 = 1 - \Pr(g|\text{monitor,shirk}) \\
\Pr(s|\text{shirk,work}) &= \tfrac{1}{2} = \Pr(s|\text{shirk,shirk})
\end{aligned}$$

for every $s \in S$. In other 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 that enters players' utility linearly (with unit constant marginal utility).

Below we consider two polar cases. In the first case, signal realizations are publicly verifiable, whereas in the second case only Robinson can verify signals.

2.1 Public Monitoring

Clearly, it's impossible to implement the efficient strategy profile (shirk,work), since no signal-contingent contractual payment could compensate Friday more when working than shirking, each signal realization carrying the same probability regardless of Friday's effort. However, we can get arbitrarily close. Let's try to implement the correlated strategy[5]

$$\sigma[(\text{monitor,work})] + (1-\sigma)[(\text{shirk,work})]$$

for any positive probability $\sigma > 0$. The best strategy for the team is to make σ as small as possible. This can be attained with the following incentive contracts, which will depend on the mediator's recommendations.

When Robinson is recommended to monitor, let $\zeta_f(s)$ denote the contractual money payment accruing to Friday if the team's signal realization is s, and let $\zeta_r(s)$ denote Robinson's associated contingent contractual payment. To simplify notation, let

$$\overline{\zeta}_i \ := \ \tfrac{1}{2}(\zeta_i(g) + \zeta_i(b))$$

be the average payment to player i in case Robinson shirks when he was told to monitor. Similarly, let $\xi_i : S \to \mathbb{R}$ be the players' contractual payments when Robinson is recommended to shirk, with $\overline{\xi}_i$ denoting average payments as before.

Implementing σ requires satisfaction of two incentive constraints, one for each player. For Robinson, it is given by

$$2 + \zeta_r(g) \ \geq \ 3 + \overline{\zeta}_r$$

when it is suggested to him that he monitors. (When Robinson is recommended to shirk, we may as well pay him $\xi_r(s) = 0$ for any s, since then he will willingly shirk.) As for Friday, his incentive constraint looks like

$$-1 + \sigma\zeta_f(g) + (1-\sigma)\overline{\xi}_f \ \geq \ \sigma\zeta_f(b) + (1-\sigma)\overline{\xi}_f.$$

Simplifying, we obtain:

$$\begin{aligned} \zeta_r(g) - \overline{\zeta}_r \ &\geq \ 1 \\ \sigma[\zeta_f(g) - \zeta_f(b)] \ &\geq \ 1. \end{aligned}$$

[5] As 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.

The first equation requires that Robinson's monitoring cost be outweighed by his contractual gain associated with monitoring as opposed to shirking. The second inequality states that for Friday to prefer working over shirking, contractual payments for him working net of effort cost must exceed (in expectation, relative to Robinson's probability of monitoring) his payment for shirking.

There clearly exist payments that implement any $\sigma > 0$ (simply make $\zeta_i(b)$ extremely low and $\zeta_i(g)$ sufficiently high), but there is no contract that implements $\sigma = 0$ (since Friday's incentive constraint would necessarily be violated). In particular, consider the following incentive contracts, which implement any $\sigma > 0$.

$$\xi_i \equiv 0, \qquad \zeta_r(g) = 2, \qquad \zeta_f(g) = 1/\sigma, \qquad \zeta_i(b) = 0.$$

This contractual arrangement employs nonnegative payments for every recommended action and signal realization. It may be viewed as a version of Holmström's (1982) *team punishments* approach to creating incentives for overcoming moral hazard in teams. It appears that both Robinson and Friday are "agents," in that the structure of their contracts is identically in line with the paradigm of team punishments: both team members are rewarded if and only if the realized signal is "good news."

2.2 Private Monitoring

Consider next the case where signal realizations are only observed in private by Robinson, and that such observations are not verifiable. Therefore, if the principal is to write signal-contingent contracts, he must solicit the realizations from Robinson who may in principle misreport them.

Notice first of all that under the previous contractual arrangement neither Robinson has the incentive to monitor nor does Friday have the incentive to work. The contract fails to be incentive compatible because now Robinson never wants to be honest. Indeed, the best strategy for him is not only to shirk, but furthermore to always report g. If he does so then he can guarantee the positive payment $\zeta_r(g) = 2$ in addition to saving himself the cost of monitoring effort. Therefore, team punishments fail to provide the right incentives for Robinson. In turn, given that Robinson's signal report to the principal is always g, Friday's contractual payments are $\zeta_f(g) = 1/\sigma$ independently of whether or not he works, rendering shirk a dominating strategy. It follows that only (shirk,shirk) is implementable with the previous contracts.

One way to convince Robinson to report honestly would be to make his contractual payments independent of reported signals. In this case Robinson would be indifferent between every reporting strategy, and in particular would be happy to report truthfully. However, if every reporting strategy yielded the same payoff then Robinson would find no incentive to incur any monitoring effort whatsoever.

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 somehow the principal and Friday could correlate their play without Robinson knowing about it, then it might be possible to "cross-check" Robinson's report, thereby "monitoring the monitor."

The following correlated strategy will be implemented:

- monitoring is recommended to Robinson with probability σ (and shirking with probability $1-\sigma$) as well as honest reporting,
- working is independently recommended to Friday with probability μ (and shirking with probability $1-\mu$), and
- the principal correlates his contractual strategy with players' recommendations, where ζ denotes the contract if (monitor,work) is recommended, and ξ denotes the contract if (monitor,shirk) is recommended (if Robinson is recommended to shirk then all contractual payments are assumed to equal zero).

Friday's incentive constraint in case the principal recommends him to work is

$$-1+\sigma\zeta_f(g) \;\geq\; \sigma\zeta_f(b).$$

Notice that this is identical to Friday's incentive constraint when he's asked to work in the previous example. (If Robinson is recommended to shirk then Friday gets nothing.) If Friday is recommended to shirk, then simply paying him $\xi_f(g)=\xi_f(b)=0$ makes the recommendation incentive compatible.

If Robinson is recommended by the principal to monitor then his payoff when he's honest and obedient is given by

$$\mu(2+\zeta_r(g))+(1-\mu)(-1+\xi_r(b)).$$

Robinson's dishonest but obedient payoff is given by

$$\mu(2+\max\{\zeta_r(b),\zeta_r(g)\})+(1-\mu)(-1+\max\{\xi_r(b),\xi_r(g)\}).$$

Finally, Robinson's dishonest and disobedient payoff is given by

$$\max\left\{\mu(3+\zeta_r(b))+(1-\mu)\xi_r(b),\mu(3+\zeta_r(g))+(1-\mu)\xi_r(g)\right\}.$$

Clearly, Robinson is willing to be honest if $\zeta_r(g)=\zeta_r(b)$ and $\xi_r(g)=\xi_r(b)$. But then, as was suggested earlier, there's no reason for him to be obedient.

Another way to make Robinson honest is to give him incentives to confirm the mediator's recommendation to Friday, i.e., let $\zeta_r(b)<\zeta_r(g)$ and $\xi_r(b)>\xi_r(g)$. Specifically, let $\zeta_r(b) = \xi_r(g) = 0$. Honesty is now assured. As for obedience, we require the honest and obedient payoff to be greater than or equal to

$$\max\left\{\mu\zeta_r(g),(1-\mu)\xi_r(b)\right\}+3\mu.$$

This leads to two inequalities in the remaining contractual unknowns. Manipulating them yields the following requirements for incentive compatibility:

$$\zeta_r(g) \;\geq\; 1/\mu, \qquad \xi_r(b) \;\geq\; 1/(1-\mu).$$

The first constraint shows that for Robinson to willingly monitor, Friday must shirk with positive probability so that Robinson truly faces payoff uncertainty regarding Friday's effort. The second constraint shows that it is impossible to have Friday working with unit probability because of Robinson's incentive problem, unless we allow for unbounded conditional rewards.

The following contingent payments implement any (σ,μ) such that $0<\sigma,\mu<1$.

$$\begin{aligned} \zeta_r(g) \;&=\; 1/\mu, \qquad \xi_r(b) \;=\; 1/(1-\mu), \qquad \zeta_r(b) \;=\; \xi_r(g) \;=\; 0,\\ \zeta_f(g) \;&=\; 1/\sigma, \qquad \zeta_f(b) \;=\; 0, \qquad \xi_f(b) \;=\; \xi_f(g) \;=\; 0. \end{aligned}$$

Friday's payments are the same as before, and Robinson's ex ante payment from monitoring is $\mu\zeta_r(g)+(1-\mu)\xi_r(b)=2$, as in the previous example. Perhaps the most distinguishing properties of this new arrangement with respect to the previous one are firstly that Robinson does not directly observe the principal's recommendation to Friday, and secondly that Robinson has the incentive to monitor to the extent that payments reward him for his reporting accuracy.

It will subsequently will be argued that the broad contractual structure derived here applies to any situation involving costly private monitoring. This broad structure will be referred to as "loyalty-testing" contracts, on the grounds that Robinson's report only confirms to the principal his own recommendation to Friday.

3 Public Monitoring and Metering Productivity

Our analysis of the metering problem begins with an environment labelled *public monitoring*. For an arbitrary team, we assume the existence and availability of a mediating principal[6] who makes behavioral recommendations to team members and implements contractual rewards contingent upon the realization of a publicly verifiable (and possibly noisy) signal, whose probability distribution in principle depends on team members' actions.

The main characterization result of this section is that team members can be provided with incentives to exert effort for the team with contracts that are structurally comparable to Holmström's (1982) *team punishments*. In other words, team members may be rewarded in case that signal realizations are "good" and punished if they are "bad." This arrangement applies not just to workers, but also to monitors, who at least contractually fail to be distinguished from other productive workers. Their only difference is really that monitoring is not directly productive. Rather, their productivity is indirect in that their function is to expand the team's contractual possibilities by providing incentives to team members.

That monitoring effort may be induced by simple incentive contracts just like a worker's compensation for productive effort refutes the claims of Alchian and Demsetz (1972). Arguably, the present model fails to capture essential aspects of monitoring, and to a large extent Section 4 aims at filling this potential void. However, as we shall see, monitoring is not substantially different from working.

We begin by laying out formal preliminaries, defining the team's problem, and presenting illustrative examples. Specifically, we formally describe a team production technology in the spirit of Makowski and Ostroy (2003), where individual members play a normal-form game and make commodity trades. We define a team's monitoring technology as well as an incentive contract in terms of such trades.

Next, we explore the meaning of our contracts a little more deeply by emphasizing their "in-kind" nature. Individuals are rewarded for their effort with commodity payments. These may be local public goods or private goods. We consider a special case where there is a quasi-linear private good called "the incentive good," in line with the Robinson and Friday examples of Section 2.

[6] See Rahman (2005b). We will use the terms "mediating principal," "principal," and "organizer" interchangeably.

Optimum contracts are then characterized with the duality theory of linear programming. We define a team's optimum contracting problem when there are only linear transfers (possibly unbounded) available, and use duality to obtain the weakest sufficient condition for approximate implementability of any team outcome. If the monitoring technology satisfies *convex independence* (a popular condition in the repeated games literature, see Section 5) then for any outcome there exist incentive contracts that implement it arbitrarily closely. We also find conditions to distinguish outcomes attainable with bounded contracts from those that are only approximately attainable.

The general structure of team punishments is confirmed implementing team outcomes with rewards and punishments, and deriving conditions on the monitoring technology to identify rewarding versus punishment information states. Finally, Holmström's contracts are shown to be a special case.

3.1 A Team Production Technology

Teams are assumed to engage in team production. Given a team consisting of finitely many individuals collected in the set $I = \{1, \ldots, n\}$, each of its members $i \in I$ has a finite set of actions available to him collected in the set A_i, with typical element a_i.

Team actions, i.e., profiles of individual actions indexed by team members, are denoted by a and belong to the product space

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

Team actions have two consequences for the team. First of all, they have repercussions on utilities. Every individual member $i \in I$ is assumed to have a utility function

$$v_i : A \times \mathbb{R}^\ell \to \mathbb{R}$$

defined on the space of team actions and that of net trades of commodities (with $\ell < \infty$).[7] Thus, if the team action is $a \in A$ and the net trade is $z \in \mathbb{R}^\ell$, the associated utility to individual i is denoted by $v_i(a, z)$.

Our first formal assumption is on the topological properties of v_i.

[7]For any $z = (z_1, \ldots, z_\ell) \in \mathbb{R}^\ell$, we adopt the convention that $z_k > 0$ represents an inflow of the kth commodity for the entire team and $z_k < 0$ represents an outflow.

Assumption 3.1. For every $i \in I$, and $a \in A$, the function $v_i(a) : \mathbb{R}^\ell \to \mathbb{R}$ defined by $v_i(a)(z) := v_i(a, z)$ is Lipschitz on its effective domain, dom $v_i(a)$, a compact, convex set containing the zero vector.

Team actions also have a direct effect on the team's trading possibilities. Every team of any type is assumed to take as given a function

$$v_0 : A \times \mathbb{R}^\ell \to \{0, -\infty\},$$

called the team's *trading possibilities indicator*, where $v_0(a, z) = 0$ means that it is technologically possible, feasible, for the team to trade z when their team action is a; the value $-\infty$ means that it is impossible. This leads us to our second assumption, that constrains the set of feasible trading possibilities.

Assumption 3.2. For every $a \in A$, dom $v_0(a) = v_0(a)^{-1}(0)$ is a compact, convex set that contains the zero trade vector.

Let the family of all utility functions relevant to the team be denoted by

$$\mathbf{v} \;=\; (v_0, v_1, \ldots, v_n).$$

Given a team action a and a trade z, I will denote the *team's utility* by $v(a, z)$ and define it by the following summation:

$$v(a, z) \;:=\; \sum_{i=1}^{n} v_i(a, z) + v_0(a, z).$$

The value $v(a, z)$ is interpreted as the team's welfare—unambiguously defined with equal welfare weights by assuming transferable utility—when the team action is a and the team's net trade is z, for any z that is feasible with respect to a. If z is not feasible with respect to a, then z lies outside the effective domain of $v(a)$.

The assumption of transferable utility is really just a normalization. When the team's problem of maximizing welfare is introduced, its non-transferable utility version is described by a family of problems indexed by $\lambda \in \Delta(I)$, with rescaled utility functions $\lambda \cdot \mathbf{v} = (v_0, \lambda_1 v_1, \ldots, \lambda_n v_n)$ and team utility given by

$$\lambda \cdot v(a, z) \;:=\; \sum_{i=1}^{n} \lambda_i v_i(a, z) + v_0(a, z),$$

for which $\lambda_i = 1/n$ corresponds to the original version with transferable utility.

3.2 The Team's Problem

For every team there is a finite set S together with a measure-valued map

$$\Pr : A \to \Delta(S)$$

called the team's *public monitoring technology*. The set S carries the interpretation of a sample space that describes the collection of contingencies upon which verifiable contracts may be enforced. Team actions lead to probabilistic realizations of such signals. Thus, if the team's action is a, then $\Pr(s|a)$ is the conditional probability that the team will publicly verify the realization s.

Some actions may lead to realizations that precisely identify the actions themselves, whereas other actions may lead to blurred signals thereof. Efficient contractual arrangements ought to exploit this fact. In order for contracts to induce productive behavior, they may require certain monitoring actions to improve upon incentives provided by the contracts themselves.

Although the contractual technology does not directly depend on the team's net trade, it is still possible for net trades to affect the verification technology indirectly. For instance, a team action may involve the use of monitoring equipment (such as a video camera) which would be impossible to implement without the purchase of such equipment. This impossibility is captured by v_0.[8]

For instance, a team's monitoring technology may have S being a singleton set or $\Pr(s|a) = \Pr(s|b)$ for every pair of team actions a, b, in which case contractual payments would be useless. On the other hand, *transparent* teams (see Makowski and Ostroy, 2003) may be described by $S = A$ together with $\Pr(s|a) = 1$ if $s = a$ and zero otherwise, in which case contractual payments would completely overcome incentive problems. Transparent teams will be discussed in more detail below.

By definition, an *incentive contract* is a map $\zeta : S \to \mathbb{R}^\ell$. The principal's action space, A_0, is the space of all such incentive contracts. The timing interpretation of these strategies is that first the principal mediates by recommending a team action and committing to a contract ζ, then active players play some a, after which a signal s realizes; finally, the principal trades $\zeta(s)$.

[8] In order to relate the current model to existing literature on moral hazard, we may assume that v_0 depends on s alone or also on a.

Given a team action a, every team member $i \in I$ has an augmented utility function $v_i(a) : A_0 \to \mathbb{R}$ over incentive contracts denoted and defined by

$$v_i(a, \zeta) \; := \; \sum_{s \in S} v_i(a, \zeta(s)) \Pr(s|a).$$

Contracts provide incentives to individuals by effectively changing team members' utility functions over team actions. The team's utility function $v(a) : A_0 \to \mathbb{R}$ from a contract ζ is also augmented to

$$v(a, \zeta) \; = \; \sum_{i=1}^{n} v_i(a, \zeta) + v_0(a, \zeta).$$

We may now write down the team's problem. For any net trade $z \in \mathbb{R}^\ell$, let[9]

$$\begin{aligned}
\langle \mathbf{v} | \Pr \rangle(z) \; := \; \sup_{\sigma \in \Delta} \quad & \sum_{(a,\zeta)} \sigma(a, \zeta) v(a, \zeta) \quad \text{s.t.} \\
& \sum_{(a_{-i},\zeta)} \sigma(a, \zeta)[v_i(b_i, a_{-i}, \zeta) - v_i(a, \zeta)] \; \leq \; 0 \\
& \sum_{(a,\zeta,s)} \sigma(a, \zeta) \Pr(s|a) \zeta(s) \; = \; z.
\end{aligned}$$

The team's problem at a given net trade z is to find a welfare-maximizing correlated strategy for the game involving the principal's contractual strategies subject to it being a correlated equilibrium and the principal's strategy agreeing with z on average.

The principal's strategy consists of signal-contingent trades, and a deviation by a player carries with it trading consequences inasmuch as the deviation might influence signal-realization probabilities.

The set of feasible solutions to the team's problem is clearly a closed and convex set, since the set of correlated equilibria is closed and convex. The following proposition is now immediate.

Proposition 3.3. *The team utility function $\langle \mathbf{v} | \Pr \rangle$ is concave in z.*

To illustrate the team's problem, consider the following example, which is a version of the principal-agent problem with output being neither random nor signal-contingent to begin with, although it optimally turns out to be random in order to overcome incentive constraints.

[9] The supremum is defined with respect to Δ, the set of regular, Borel probability measures on $A \times A_0$; for any $f : A \times A_0 \to \mathbb{R}$, the sum $\sum_{(a,\zeta)} f(a, \zeta)\sigma(a, \zeta)$ is shorthand for $\int_{A \times A_0} f(a, \zeta) d\sigma(a, \zeta)$.

Example 3.4. Suppose that $\ell = 1$ and that there is only one active player $i \in I$, the "agent." He is assumed to have two available actions: $A = \{\text{work,shirk}\}$. Utility functions are restricted to have the unit interval $[0,1]$ as their effective domain, and assume the following functional forms:

$$\begin{aligned} v_i(\text{work}, z) &= \min\{2z, \tfrac{1}{2}z + \tfrac{1}{2}\} = v_i(\text{shirk}, z) - \tfrac{1}{3} \\ v(\text{work}, z) &= z = v(\text{shirk}, z) + \tfrac{1}{2} \end{aligned}$$

The team's contractual technology is given by the set $S = \{\text{good,bad}\}$ of possible news together with the probabilities

	good	bad
work	3/4	1/4
shirk	1/3	2/3

where each entry denotes the associated value of Pr.

The following incentive contract induces the agent to work: let $\zeta(\text{good}) = 1$ and $\zeta(\text{bad}) = 0$. At $z = \frac{3}{4}$, the correlated strategy $\sigma(\text{work}, \zeta) = 1$ is a feasible solution of the team's problem, since the agent's incentive constraint $\frac{3}{4} \geq \frac{1}{3} + \frac{1}{3} = \frac{2}{3}$ is satisfied. The team's utility therefrom is calculated to be $\frac{3}{4}$. No other incentive contract can improve upon this level of utility for the team, therefore $\langle \mathbf{v} | \Pr \rangle(\frac{3}{4}) = \frac{3}{4}$.

To further illustrate the team's problem, let us derive Makowski and Ostroy's (2003) transparent teams (i.e., with access to binding contracts) as a special case.[10]

Example 3.5. Let $S = A$, with $\Pr(s|a) = 1$ if $s = a$ and zero otherwise. Then

$$\begin{aligned} \langle \mathbf{v} | \Pr \rangle(z) = \sup \quad & \sum_{(a,\zeta)} \sigma(a,\zeta) v(a, \zeta(a)) \quad \text{s.t.} \\ & \sum_{(a_{-i},\zeta)} \sigma(a,\zeta)[v_i(b_i, a_{-i}, \zeta(b_i, a_{-i})) - v_i(a, \zeta(a))] \leq 0 \\ & \sum_{(a,\zeta)} \sigma(a,\zeta)\zeta(a) = z. \end{aligned}$$

Clearly, ignoring the incentive constraints above, the team's problem simplifies to $(\text{conc } \mathbf{v})(z) := \sup\{\sum_{(a,\hat{z})} \sigma(a,\hat{z}) v(a,\hat{z}) : \sum_{(a,\hat{z})} \sigma(a,\hat{z})\hat{z} = z\}$. Perhaps less clear is that with incentive constraints, the value of the team's problem still amounts to conc $\mathbf{v}$. In other words, we have the following equality of functions.

[10] Equivalently, the behavior of team members is freely observable and contractible.

Proposition 3.6. *For any transparent team,* $\langle \mathbf{v} | \Pr \rangle \equiv \operatorname{conc} \mathbf{v}$.

For proof, we resort to "as-if binding" contracts. Let us say that the principal is not constrained to feasible trading strategies. In particular, he may strike action-contingent trades that lie outside dom $v(a)$. By definition, an *as-if binding contract* for a given $b \in A$ is a transparent incentive contract $\zeta_b : A \to \mathbb{R}^\ell$ such that

$$\zeta_b(a) \in \bigcup_{i=1}^{n} \operatorname{dom} v_i(a)$$

if and only if $a = b$. By Assumption 3.1, $\bigcup \operatorname{dom} v_i(a)$ is non-empty and compact, so its complement is also non-empty, hence as-if binding contracts are well-defined. Given $b \in A$, an as-if binding contract ζ_b implements b and relaxes all players' incentive constraints, since if any player i contemplated a deviation from b_i to a_i then he would obtain a utility of $-\infty$ from $\zeta_b(a_i, b_{-i})$. Therefore, by resorting to such incentive contracts if necessary, Proposition 3.6 now follows.

3.3 Linear Transfers with the Incentive Good

The team's problem so far has taken a prescribed trade of commodities z as given, and allocated it contractually with some ζ to maximize welfare subject to an average resource constraint, viewing z and ζ as a public good for the team, i.e., every individual consumed the same ζ. This includes the case of contracts denominated in private goods if we view a team's allocation of private goods as a local public good. For instance, with only private goods in the economy, one may augment the domain of $\langle \mathbf{v} | \Pr \rangle$ to $\mathbb{R}^{\ell \times n}$: the allocation of private goods to the team's members.

On the other hand, if allocating private goods to one individual is perfectly substitutable with allocating it to any other individual, the team's problem may be specialized to incorporate the private goods allocation problem.

Specifically, suppose that there is only a so-called *incentive good*[11] that enters linearly every team member's utility as follows:

$$v_i(a, z_i) = v_i(a) + z_i.$$

[11] It is not called "money" because in a general equilibrium setting with transferable utility this might not lead to a well-defined model. See Rahman (2005b).

The incentive good contradicts Assumption 3.1, since dom $v_i(a) = (-\infty, \infty)$ is not compact. An *incentive contract* is now a map $\zeta : I \times A \times S \to \mathbb{R}$, since—every player being risk neutral with respect to z_i—a correlated strategy σ need not randomize with respect to ζ above and beyond the organizer's recommendation, a. Therefore, the team's problem may be formulated as:[12]

$$\langle \mathbf{v} | \Pr \rangle(z) := \sup_{\sigma \in \Delta, \zeta} \sum_{(i,a,s)} \sigma(a)[v_i(a) + \zeta_i(a,s) \Pr(s|a)] \quad \text{s.t.}$$

$$\sum_{(a_{-i},s)} \sigma(a)[(v_i(b_i, a_{-i}) + \zeta_i(a,s)) \Pr(s|b_i, a_{-i}) - (v_i(a) + \zeta_i(a,s)) \Pr(s|a)] \leq 0$$

$$\sum_{(i,a,s)} \sigma(a) \Pr(s|a) \zeta_i(a,s) = z.$$

Non-compactness of dom $v_i(a)$ leads to phenomena like non-existence of a solution to the problem above. For instance, in Section 2.1, the profile (shirk,work) was not attainable, yet it could be approached arbitrarily closely. Indeed, recall that from the arrangement proposed in Section 2.1, expected contractual payments to team members given $\sigma > 0$ amount to

$$z = \sigma(2 + 1/\sigma) + (1 - \sigma)0 = 1 + 2\sigma.$$

This suggests that $\sigma = 0$ is the cheapest way to induce Friday to work by paying him an unbounded amount in case Robinson monitors. In fact, the team's utility function over the incentive good reflects this. (Assume $\zeta_i \geq 0$, interpreted as *limited liability.*)

Claim 3.7. *The team's utility from Section 2.1 is given by*

$$\langle \mathbf{v} | \Pr \rangle(z) = \begin{cases} 3z & \text{if } 0 \leq z < 1 \\ 2 + z & \text{if } z \geq 1 \\ -\infty & \text{if } z < 0. \end{cases}$$

However, the supremum in the team's problem is never attained.

For $0 \leq z < 1$, the team's utility is calculated by having it play (shirk,shirk) without payment with probability $(1 - z)$ and (shirk,work) with "unbounded" payments with probability z. For $z \geq 1$, the team plays (shirk,work), which again is implemented by "unbounded" payments. Such "unbounded" contract is understood as the limit when σ tends to zero of the contract proposed in Section 2.1.

[12] Strictly speaking, this version of the team's problem is not a linear program. By choosing $\sigma(a)$ and $\sigma(a)\zeta_i(a,s)$ independently, it becomes linear. An optimum such that $\sigma(a)\zeta_i(a,s) \neq 0$ yet $\sigma(a) = 0$ is interpreted to mean that unbounded payments are necessary, as in Claim 3.7.

3.4 Optimum Contracts via Duality

To better understand the nonexistence result in Claim 3.7, consider the following example, which shows that although the concave extension of two functions may be well defined, the supremum may not actually be attained. Let $v_1(z) = \min\{2z, 2\}$ and $v_2(z) = z$ be defined for $z \geq 0$. It is easy to see that $\text{conc}\{v_1, v_2\}(z) = \min\{2z, 1+z\}$, yet calculating $\text{conc}\{v_1, v_2\}(z) = \sup\{\sigma v_1(z_1) + (1-\sigma)v_2(z_2) : \sigma z_1 + (1-\sigma)z_2 = z\}$, the supremum is not attained by a finite z_2 when $z > 1$. A similar effect is taking place in Robinson and Friday's attainment problem.

Using the insight of Footnote 11, we will now derive a linear program that, with unbounded contracts if necessary, characterizes a team's contractual possibilities. Consider the following *primal* problem:

$$
\begin{aligned}
\langle \mathbf{v} || \Pr \rangle(z) \ := \quad \sup_{\sigma \geq 0, \xi} \ & \sum_{(i,a)} \sigma(a) v_i(a) + \sum_{(i,a,s)} \xi_i(a,s) \Pr(s|a) \quad \text{s.t.} \\
\sum_{a_{-i}} \sigma(a)[(v_i(b_i, a_{-i}) - v_i(a)] + \sum_{(a_{-i},s)} \xi_i(a,s)[\Pr(s|b_i, a_{-i}) - \Pr(s|a)] \ & \leq \ 0 \\
\sum_{(i,a,s)} \xi_i(a,s) \Pr(s|a) \ & = \ z \\
\sum_{a \in A} \sigma(a) \ & = \ 1.
\end{aligned}
$$

The team chooses a correlated strategy $\sigma \in \Delta(A)$ and a function $\xi : I \times A \times S \to \mathbb{R}$ of (probability weighted) payments to maximize welfare, subject to the following constraints. The last constraint makes sure that σ is a probability measure, the penultimate constraint makes sure that payments add up to the team's resources of the incentive good, and finally the first family of constraints, indexed by $i \in I$ and $a_i, b_i \in A_i$, makes sure that obeying the organizer is incentive compatible.

This problem would agree with the previous one if we added the constraint that $\xi_i(a,s) = \sigma(a)\zeta_i(a,s)$ for some ζ_i. We will adopt a "converse" approach. Whenever there is no ζ_i to satisfy this added requirement, it will be interpreted that unbounded contracts are necessary. Indeed, it is easy to see that no ζ_i exists if and only if $\sigma(a) = 0$ yet $\xi_i(a,s) \neq 0$. Alternatively, to focus on finite payments while maintaining a linear programming structure, we might have included bounding constraints such as

$$-\sigma(a)\overline{z} \ \leq \ \xi_i(a,s) \ \leq \ \sigma(a)\overline{z}.$$

Later, we will show that this converges to the primal above as $\overline{z} \to +\infty$.

First, though, notice that with the Lagrangian of this primal and a little symbolic arithmetic, it is not difficult to show that the associated *dual* problem is given by the following linear program:

$$
\begin{aligned}
\langle \mathbf{v} \| \Pr\rangle(z) &= \inf_{\lambda \geq 0, \mu, \nu} \quad \mu z + \nu \quad \text{s.t.} \\
\sum_{i=1}^{n} (1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i)) v_i(a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) v_i(b_i, a_{-i}) &\leq \nu \\
(1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i)) \Pr(s|a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) \Pr(s|b_i, a_{-i}) &= \mu \Pr(s|a).
\end{aligned}
$$

The dual minimizes some "expenditure" $\nu + \mu z$ by choosing $\lambda \geq 0$, which indexes the multipliers on the primal incentive constraints, μ, which is the multiplier associated with the team's resource constraint (with respect to the incentive good), and ν, which is the multiplier associated with the the probability constraint on σ.

The first family of dual constraints, indexed by $a \in A$, arises from the primal when σ is interpreted as a multiplier. The second family, indexed by $i \in I$, $a \in A$, and $s \in S$, arises when ξ is interpreted as a multiplier. (That the dual value equals the primal value follows from the Fundamental Theorem of Linear Programming.)

Almost immediately, the dual reveals the following result.

Lemma 3.8. *If (λ, μ, ν) solves the dual then $\mu = 1$. Therefore,*

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

for every $i \in I$, $a \in A$, and $s \in S$, at any dual solution.

Proof. Adding the second family of dual constraints with respect to $s \in S$ for each (i, a), and noting that $\sum_s \Pr(s|a) = 1$ for every $a \in A$, it follows that

$$
\begin{aligned}
\left[1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i)\right] \sum_{s \in S} \Pr(s|a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) \sum_{s \in S} \Pr(s|b_i, a_{-i}) &= \mu \sum_{s \in S} \Pr(s|a) \\
\Rightarrow \quad 1 = 1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) &= \mu,
\end{aligned}
$$

proving the first claim. The second claim follows by substituting $\mu = 1$ into the dual constraints and subtracting $\Pr(s|a)$ from both sides. □

Lemma 3.8 motivates our next definition, which provides sufficient conditions for any correlated strategy to be implementable (with possibly unbounded payments).

Definition 3.9. A public monitoring technology Pr satisfies *convex independence* if for every player $i \in I$ and individual action $a_i \in A_i$,

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

where $\Pr[a_i] : A_{-i} \to \Delta(S)$ is the family of conditional probability vectors for s given a indexed by a_{-i} and conv stands for convex hull.[13]

Convex independence suffices for any team to maximize welfare and overcome its incentive constraints. It is also the *weakest* sufficient condition: without it, there exist preferences for team members such that incentive constraints will bind.

Theorem 3.10. *The public monitoring technology* Pr *satisfies convex independence if and only if*

$$\langle \mathbf{v} || \Pr \rangle(0) \;=\; \max\{v(a) : a \in A\}$$

for every $\mathbf{v} \in \mathbb{R}^{nA}$*, where* $v := \sum_i v_i$ *and* $n > 1$*.*

Proof. For sufficiency, if Pr satisfies convex independence then

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

for every (i, a) and $\lambda \geq 0$ implies that $\lambda_i(a_i, b_i) = 0$ for all (i, a_i, b_i) with $a_i \neq b_i$. By Lemma 3.8 this must hold at any dual solution. Therefore, looking at the first family of dual constraints, it follows that any dual solution must satisfy

$$\sum_{i \in I} v_i(a) \leq \nu$$

for every $a \in A$. Minimizing ν subject to these constraints yields the claimed equality.

For necessity, if convex independence fails then there exists nonzero λ that satisfies the equalities of Lemma 3.8. Choose λ to be nonzero only for one $(j, \widehat{a}_j)$, and pick $\mathbf{v}$ as follows. For any a_{-j}, the utility to each player depending on whether or not j plays $\widehat{a}_j$ is given by (first is j then anyone else):

a_j	$\widehat{a}_j$
1, 0	0, 2

[13]This is weaker than Definition 5.1 in Fudenberg et al. (1994), imposing *linear* independence.

For any a with $a_j \neq \widehat{a}_j$, the first dual constraint becomes $1 \leq \nu$, since $\lambda = 0$ there. At $\widehat{a}_j$, the constraint becomes $2(n-1) - \sum_{b_j} \lambda_j(\widehat{a}_j, b_j) \leq \nu$. Since $\sum \lambda > 0$, there is a feasible dual solution with $\nu < 2(n-1) = \max\{v(a)\}$, as required. □

Intuitively, if Pr satisfies convex independence then a player's action may be detected, in that there are strategy profiles that distinguish it probabilistically from any deviation, or combination thereof. From the dual, convex independence is equivalent to the statement that there are contractual payments that relax a team's incentive constraints, which naturally leads to Theorem 3.10.

For instance, the monitoring technology in Section 2.1 satisfies convex independence. In that example, the profile (monitor,work) was attainable with bounded contracts, yet the profile (shirk,work) was not. Next, we will describe what it is about Pr that makes some profiles attainable but not others. To this end, a correlated strategy σ is called *approachable* if there is a feasible solution to the primal (σ, ξ) for some ξ. The correlated strategy σ is called *attainable* if there is a feasible solution to the primal (σ, ξ) such that $\sigma(a) = 0$ implies $\xi_i(a, s) = 0$ for every (i, s).

Corollary 3.11. *If* Pr *satisfies convex independence then every correlated strategy is approachable.*

Proof. Replace the primal objective with just $\sum \xi_i(a, s) \Pr(s|a)$, and, without loss, suppose $z = 0$. This problem characterizes approachable correlated strategies, its value equals zero at an optimum. The dual of this problem is similar to the original dual except that the constraints associated with σ now look like

$$\sum_{i=1}^{n} \sum_{b_i \in A_i} \lambda_i(a_i, b_i) v_i(a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) v_i(b_i, a_{-i}) \ \leq \ \nu.$$

Since Pr satisfies convex independence, $\lambda_i(a_i, b_i) = 0$ for all (i, a_i, b_i) with $a_i \neq b_i$ at any dual solution, so the constraints become $0 \leq \nu$ for every $a \in A$, implying that $\nu = 0$ at a dual solution. Since any $\sigma \in \Delta(A)$ satisfies complementary slackness with respect to these constraints, the result follows. □

Definition 3.12. Pr satisfies *convex independence at* $\sigma \in \Delta(A)$ if for every $i \in I$ and $a_i \in A_i$ such that $\sigma(a_i) > 0$,

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

where $\sigma[a_i] \Pr[b_i](a_{-i}) := \sigma(a) \Pr(b_i, a_{-i})$ and $\Pr(b_i, a_{-i}) \in \Delta(S)$.

Clearly, Pr satisfies convex independence if and only if it satisfies convex independence at σ defined by $\sigma(a) = 1/|A|$ for every $a \in A$. Moreover, it is easy to see that Pr satisfies convex independence at any σ with full support (i.e., with $\sigma(a) > 0$ for every a).[14] This strategy-specific version of convex independence implies attainability, i.e., there are bounded payments ζ that make a correlated strategy incentive compatible.

Proposition 3.13. *If* Pr *satisfies convex independence at σ then σ is attainable.*

Proof. Consider the following problem, which is a version of that in Section 3.3:

$$\begin{aligned} \langle \mathbf{v}|\,\Pr,\sigma\rangle \;:=\; \sup_{\zeta} \sum_{(i,a,s)} \sigma(a)\zeta_i(a,s)\Pr(s|a) \quad \text{s.t.} \\ \sum_{(a_{-i},s)} \sigma(a)[(v_i(b_i,a_{-i})+\zeta_i(a,s))\Pr(s|b_i,a_{-i}) - (v_i(a)+\zeta_i(a,s))\Pr(s|a)] \;\le\; 0 \\ \sum_{(i,a,s)} \sigma(a)\Pr(s|a)\zeta_i(a,s) \;=\; 0. \end{aligned}$$

This linear program takes a correlated strategy σ as given and finds payments ζ that make it incentive compatible. If there exist such payments, then the value of the primal equals zero, otherwise it equals $-\infty$. The dual of this problem is given by:

$$\begin{aligned} \langle \mathbf{v}|\,\Pr,\sigma\rangle \;=\; \inf_{\lambda\ge 0,\mu} \sum_{(i,a_i,b_i)} \lambda_i(a_i,b_i) \sum_{(a_{-i},s)} \sigma(a)[v_i(a)-v_i(b_i,a_{-i})] \quad \text{s.t.} \\ \sigma(a)\Pr(s|a)(1-\mu) \;=\; \sigma(a)\sum_{b_i\in A_i} \lambda_i(a_i,b_i)[\Pr(s|b_i,a_{-i})-\Pr(s|a)], \end{aligned}$$

where the constraints range across all (i,a,s). Adding the constraints with respect to $s \in S$, it follows that $\mu = 1$ at any dual solution. Now, if Pr satisfies convex independence at σ then the constraints are satisfied only if $\lambda_i(a_i,b_i) = 0$ for every a_i such that $\sigma(a_i) > 0$ and $b_i \neq a_i$. (Let λ be anything nonnegative elsewhere.) Not only is this feasible, but it solves the dual with value equal to zero. By the Fundamental Theorem of Linear Programming, there is a primal solution with the same value. □

In Section 2.1, Pr satisfies convex independence at (monitor,work), and as such is attainable. However, this fails at (shirk,work), so it is only approachable. Incidentally, the value of the dual in the proof of Proposition 3.13 is $-\infty$ at (shirk,work), implying that the value of the primal is also $-\infty$, i.e., there is no feasible ζ that implements it. With transparent teams (see Example 3.5), every correlated strategy is attainable, since Pr exhibits convex independence everywhere.

[14]For proof, see Lemma 4.16.

On the other hand, if $\Pr(s|a) = \Pr(s)$ for every $a \in A$ (so actions do not change the signal distribution) then any nonnegative λ is feasible (for the three duals discussed). This characterizes correlated equilibrium (see also Myerson, 1997): a given correlated strategy σ is a correlated equilibrium if and only if

$$\sum_{(i,a_i,b_i)} \lambda_i(a_i, b_i) \sum_{(a_{-i},s)} \sigma(a)[v_i(a) - v_i(b_i, a_{-i})] \; = \; 0$$

for every $\lambda \geq 0$. This might be interpreted as an extreme case of convex *dependence*.

We conclude this section by justifying the notion of approachability as the limit of attainable contracts. To illustrate (see also Section 2.1), recall that convex independence implies convex independence at every correlated strategy σ with full support, making σ attainable. Consider an approachable correlated strategy σ_0 that is not attainable, so for any feasible (σ_0, ξ) there exists (i, a, s) with $\sigma_0(a) = 0$ yet $\xi_i(a, s) > 0$. Given a sequence $\{\sigma_m\}$ of correlated strategies with full support converging to σ_0, there exist payments ζ^m that implement σ_m such that $\sigma_m\zeta^m$ converges to ξ. This follows easily by looking at the primal in the proof of Proposition 3.13: since the feasible set is defined by finitely many weak linear inequalities in a finite-dimensional space, it is closed, so every convergent sequence has a limit. With bounded utilities, without loss $\{\sigma_m\zeta^m\}$ lies in some compact interval, hence has a convergent subsequence.

To argue this claim with a general monitoring technology, we study a family of bounded versions of $\langle \mathbf{v} || \Pr \rangle$. For simplicity, suppose that $z = 0$. By definition, the *bounded primal* is the linear program that changes $\langle \mathbf{v} || \Pr \rangle(0)$ by imposing that $-\sigma(a)\overline{z} \leq \xi_i(a, s) \leq \sigma(a)\overline{z}$ for every (i, a, s), where $\overline{z} \geq 0$ is given. Denote by $\langle \mathbf{v} || \Pr, \overline{z} \rangle$ the optimum value of this perturbed primal.

Since the bounded primal is a more constrained version of the original primal, the following result obtains immediately.

Lemma 3.14. $\langle \mathbf{v} || \Pr, \overline{z} \rangle \leq \langle \mathbf{v} || \Pr \rangle(0)$ *for every* $\overline{z}$.

Next, we will derive the *bounded dual*, i.e., the dual associated with the bounded primal. We will use this linear program to show that indeed $\langle \mathbf{v} || \Pr, \overline{z} \rangle \to \langle \mathbf{v} || \Pr \rangle(0)$ as $\overline{z} \to \infty$, as well as that contractual payments may optimally take the form of Holmström's (1982) team punishments and rewards.

Letting φ be the multipliers on the constraints $-\sigma(a)\overline{z} \leq \xi_i(a, s)$ and ψ be the multipliers on $\xi_i(a, s) \leq \sigma(a)\overline{z}$, the bounded dual is given by the linear program below.

$$\langle \mathbf{v} || \Pr, \overline{z} \rangle \;=\; \inf_{\lambda,\varphi,\psi \geq 0,\mu,\nu} \quad \nu \quad \text{s.t.}$$

$$\sum_{i=1}^{n}(1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i))v_i(a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i)v_i(b_i, a_{-i}) + \overline{z}\sum_{s \in S}(\varphi_i(s|a) + \psi_i(s|a)) \;\leq\; \nu$$

$$(1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i)) \Pr(s|a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i) \Pr(s|b_i, a_{-i})$$

$$= \; \mu \Pr(s|a) + \psi_i(s|a) - \varphi_i(s|a).$$

The only difference between the bounded dual and the original dual is that now there are some extra terms in the dual constraints involving φ and ψ.

Proposition 3.15. $\langle \mathbf{v} || \Pr, \overline{z} \rangle \to \langle \mathbf{v} || \Pr \rangle(0)$ *as* $\overline{z} \to \infty$.

Proof. Let $\pi_i(a|\overline{z}) = \sum_s \varphi_i(s|a) + \psi_i(s|a) \geq 0$ at a dual solution for $\overline{z}$. We will show that $\pi_i(a|\overline{z}) \to 0$ for every (i, a). If not, then $\overline{z}\sum_i \pi_i(a|\overline{z}) \to \infty$ for some a, which, by the first family of dual constraints requiring that

$$\sum_{i=1}^{n}(1 + \sum_{b_i \in A_i} \lambda_i(a_i, b_i))v_i(a) - \sum_{b_i \in A_i} \lambda_i(a_i, b_i)v_i(b_i, a_{-i}) + \overline{z}\pi_i(a|\overline{z}) \;\leq\; \nu,$$

implies $\nu \to \infty$, contradicting Lemma 3.14. Therefore, $\varphi, \psi \to 0$, since $\varphi, \psi \geq 0$, so the dual solutions converge to a dual solution of the original primal, as required. □

Adding the second dual constraints with respect to $s \in S$, it follows that any dual solution must satisfy, for every (i, a),

$$1 \;=\; \mu + \sum_{s \in S} \psi_i(s|a) - \varphi_i(s|a).$$

(Note also that $\sum_s \psi_i(s|a) - \varphi_i(s|a)$ does not depend on (i, a).) By Proposition 3.15, $\sum_s \psi_i(s|a) - \varphi_i(s|a) \to 0$ as $\overline{z} \to \infty$. Therefore, $\mu \to 1$. For any (i, a), $S_\varphi(i, a) := \{s : \varphi_i(s|a) > 0\}$ is the set of *rewarding signals* and $S_\psi(i, a) := \{s : \psi_i(s|a) > 0\}$ is the set of *punishing signals*. Clearly, $S_\varphi(i, a) \cap S_\psi(i, a) = \emptyset$, since both constraints for which φ and ψ are multipliers cannot bind simultaneously.

To obtain Holmström's (1982) contracts, assume a is efficient (but not an equilibrium) and there exist $\overline{s}, \underline{s} \in S$ with $\Pr(\overline{s}|a) > \Pr(\overline{s}|b_i, a_{-i})$ and $\Pr(\underline{s}|a) < \Pr(\underline{s}|b_i, a_{-i})$ for every (i, b_i). The second dual constraints are now satisfied only if either $\psi_i(\overline{s}|a) > 0$ for every i, or $\varphi_i(\underline{s}|a) > 0$ for every i, or both. By complementary slackness, either $\zeta_i(a, \overline{s}) = \overline{z}$ for every i, or $\zeta_i(a, \underline{s}) = -\overline{z}$ for every i, or both: team members are rewarded and punished together.

4 Private Monitoring and Metering Rewards

This section builds on the previous one by sudying environments with *private monitoring*. We assume that signal realizations are private information that is reported by individual members to a principal, who subsequently implements report-contingent contracts. Therefore, not only must incentives be provided for individuals to obey the principal's recommendations to, say, work, but additionally team members must have incentives to truthfully report their privately observed monitoring signals.

Private monitoring has important implications for a team over and above its public counterpart. Standard team punishments are no longer incentive compatible, since monitors prefer not to report truthfully under them, and therefore find no reason to make any monitoring effort. With monitors failing to monitor, workers have no reason to work, since their rewards cease to be aligned with effort. Furthermore, to reward monitors with the same payment regardless of their reports, despite enabling truthful reporting, fails to provide incentives for any monitoring effort.

It is possible to induce both truthful reporting and monitoring effort by having his contractual rewards depend on other players' recommendations (and not revealing such recommendations to him). A monitor now has the incentive to exert effort in order to increase his chances of reporting profitably. Truthful reporting obtains by effectively asking monitors to confirm the mediator's recommendations to others. Once monitors' effort and truthful reports are enforced, all other workers may be subject to the usual team punishments.

We begin this section by formally defining a team's private monitoring technology, followed by the development of a team's *metering problem* of finding an efficient organization. By assumption, signal realizations do not take place until after actions have been adopted; this leads us to consider *sequentially rational* reporting strategies. An organization will be understood as a *sequential communication equilibrium* in the spirit of Myerson (1986).

The question of how to choose a monitor is then explored. A potential trade-off between monitoring ability and tastes over the team's trades is identified when the team does not have access to linear transfers. Finally, we focus on teams whose only contractual medium is linear transfers. Using duality, we replicate the results of Section 3.4 to private monitoring environments by amending only slightly our notion of *convex independence*.

4.1 Private Monitoring Technology

As usual, the team plays a normal-form game denoted just as before. We begin with a family $\{S_i : i \in I \cup \{0\}\}$ such that S_i is a finite set of *private signals* observable only by individual member i. The set S_0 may be interpreted as the signals observed by the principal, or as publicly observed signals. Denote by

$$S \; := \; \prod_{i=0}^{n} S_i$$

the product space of all observable signals. The team's *private monitoring technology* is given by the space S together with a measure-valued map

$$\Pr \; : \; A \to \Delta(S)$$

where $\Pr(s|a)$ stands for the conditional probability that s will be privately observed by the players given that the team adopts action profile a.

Assumption 4.1. For every $s \in S$ there exists $a \in A$ such that $\Pr(s|a) > 0$.

The timing of team members' interaction runs as follows. Firstly, players agree on a correlated strategy σ of the game augmented by the principal's contractual strategies, which in this case is the space of incentive contracts that depend on *reported* signals. Once actions have been adopted, the players report their private information (given by an element of their personal signal space) to the principal, who finally implements contracts according to the previously agreed-upon correlated strategy σ.

An incentive contract is still a map $\zeta : S \to \mathbb{R}^\ell$. Assuming that every player will choose to report truthfully, we denote the utility of player i (before actually observing his signal) from a contract ζ when the team's action profile is a by

$$v_i(a, \zeta) \; = \; \sum_{s \in S} v_i(a, \zeta(s)) \Pr(s|a).$$

Of course, Mr. i may choose to lie about his privately observed signal. A *reporting strategy* is a map $\rho_i : S_i \to S_i$ where $\rho_i(s_i)$ is the reported signal when Mr. i privately observes s_i. When all other players are truthful, the utility to i from ρ_i equals

$$v_i(a, \zeta|\rho_i) \; = \; \sum_{s \in S} v_i(a, \zeta(\rho_i(s_i), s_{-i})) \Pr(s|a).$$

A *truthful reporting strategy* is the identity map $\tau_i : S_i \to S_i$ with $\tau_i(s_i) = s_i$. Thus, $v_i(a, \zeta|\tau_i) = v_i(a, \zeta)$.

4.2 Sequentially Rational Reporting

The environment above is a two-stage game where the first stage involves no private information together with actions in A, and the second stage consists of private information S being revealed to players as well as possibly being solicited by the mediator so as to ultimately implement the team's report-contingent contractual payments in A_0. In this environment, it is conceivable that after some deviation by a player it might not be optimal for some other player to truthfully report the deviation. In defining the team's problem, such concern will now be addressed.

For instance, a monitor may prefer not to report that a worker shirked when he actually did shirk. If prescribed behavior involved the worker working with unit probability then the incentive constraints in the team's problem must not fail to include the intuitive requirement that the monitor ought to prefer to report shirking if it took place, even if—in equilibrium—it is never actually required of him to report shirking. To illustrate, consider the following example.

Example 4.2. Suppose that $\ell = 1$ and that the team consists of two players, 1 and 2. Player 1 has two actions and his signal space is trivial (a singleton set). Player 2 has a trivial action space and observes two signals. $A = \{w, s\}$ is the set of actions of player 1 and $S = \{g, b\}$ the set of signals of player 2. Signal probabilities are $\Pr(g|w) = 1 = \Pr(b|s)$ and $\Pr(b|w) = 0 = \Pr(g|s)$.

Utility functions have effective domain $[0, 1]$, and are defined by

$$\begin{aligned} v_1(w, z) &= z = v_1(s, z) - \tfrac{1}{2} \\ v_2(w, z) &= z = v_2(s, z) + 1. \end{aligned}$$

The team's utility is the sum of each player's, so $v(w, z) = 2z$ and $v(s, z) = 2z - \frac{1}{2}$.

The best action for this team is for player 1 to play w. For it to be implementable, we require an incentive contract ζ, possibly depending on player 2's reports, that leads player 1 to prefer playing w over s and leads player 2 to prefer honest reporting.

Formally, we require the satisfaction of the following three incentive constraints:

$$\begin{aligned} v_1(w, \zeta(g)) &\geq v_1(s, \zeta(b)) \\ v_2(w, \zeta(g)) &\geq v_2(w, \zeta(b)) \\ v_2(s, \zeta(b)) &\geq v_2(s, \zeta(g)). \end{aligned}$$

The first constraint requires player 1 to prefer working over shirking. The second constraint requires player 2 to prefer reporting that player 1 worked when he actually worked, and the third constraint requires player 2 to prefer reporting that player 1 shirked if he ever does shirk.

From player 1's constraint we require that $\zeta(g) \geq \zeta(b) + \frac{1}{2}$, and from player 2's first constraint we require that $\zeta(g) \geq \zeta(b)$. To satisfy the last constraint, we require that $\zeta(b) \geq \zeta(g)$, implying that $\zeta(g) = \zeta(b)$. This is inconsistent with player 1's constraint, therefore w is not implementable.

Player 2's beliefs regarding player 1's deviation were assumed to be correct in this last incentive constraint. Indeed, given his signal structure, this assumption seems reasonable. However, if player 2's signals were conditionally noisy, then the question of player 2's beliefs off the equilibrium path would have more substance, a question that leads us to reconsider sequential rationality.

What should monitors believe about their opponents' behavior upon observing signals that lead them to conclude that someone deviated? Answering this question forces us to consider *sequential rationality*, which together with correlated equilibrium in the first stage suggests Myerson's (1986) *sequential communication equilibrium* as our solution concept for this two-stage game.

Before addressing such question in general, consider two examples, which involve only changes in Pr relative to Example 4.2. Suppose first that $\Pr(g|w) = \Pr(b|w) = \frac{1}{2}$ and that $\Pr(b|s) = 1 = 1 - \Pr(g|s)$. According to this techonology, player 2's private signal amounts to a coin toss if player 1 plays w, whereas if he plays s then the signal will be b for sure. If w is implemented then both signal-realizations g and b are *on the equilibrium path*, so player 2's beliefs contingent upon each signal realization ought to be that player 1 played w. Hence, player 2 must have the incentive to report truthfully either g or b given that he believes that player 1 played w.

Now consider the "opposite" scenario, where $\Pr(g|s) = \Pr(b|s) = \frac{1}{2}$ and $\Pr(g|w) = 1 = 1 - \Pr(b|w)$. This technology assures the realization g if player 1 plays w, whereas if he chooses to play s, then player 2's signal realization amounts to a coin toss. When player 1 is meant to play w, the realization b is off the equilibrium path. Given Pr, player 2's rational beliefs concerning player 1's actions upon observing b can be nothing other than that player 1 played s.

This motivates the following approach to the team's metering problem.

First of all, the family of incentive constraints

$$\sum_{(a_{-i},\zeta,s)} \sigma(a,\zeta)[v_i(a,\zeta(s))\Pr(s|a) - v_i(b_i,a_{-i},\zeta(\rho_i(s_i),s_{-i}))\Pr(s|b_i,a_{-i})] \ \geq \ 0 \qquad (*)$$

indexed by i, a_i, b_i, and ρ_i, precludes not only disobedience (to the mediator) by a player when all other players are honest and obedient, but also dishonesty on the part of a player who chooses to deviate, since he is able to simultaneously disobey and lie about his private signal.

Assuming that other players are both honest and obedient, a given player's beliefs regarding others' behavior are not affected by his own deviation, they are just his beliefs in others' recommendations and reporting strategies. Therefore, a player's conditional beliefs on the equilibrium path are naturally given by Bayes' rule.

All that remains is to characterize players' incentives to report signals honestly *off the path of play*, noting that the incentive constraints above already consider histories where a player disobeys his own recommendation.

If a player's observation s_i has positive probability given his recommendation a_i then his beliefs about what others did ought to be given by Bayes' rule assuming that everyone was planning to be honest and obedient. If the probability of observing s_i equals zero given the recommendation a_i, then (by the previous paragraph) player i would be sure that someone had deviated. What would be reasonable beliefs for him to have upon observing such a signal? We will now answer this question.

A *history* is a pair (a_i, s_i) with $a_i \in A_i$ a recommendation to player i and $s_i \in S_i$ a privately observed signal. Call a history (a_i, s_i) *on the path of play* according to σ if

$$\Pr(a_i, s_i|\sigma) \ = \sum_{(a_{-i},\zeta,s_{-i})} \sigma(a,\zeta)\Pr(s|a) \ > \ 0,$$

and call a history (a_i, s_i) *disobedient* according to σ if

$$\sigma(a_i) \ = \sum_{(a_{-i},\zeta)} \sigma(a,\zeta) \ > \ 0$$

yet $\Pr(a_i, s_i|\sigma) = 0$. In other words, (a_i, s_i) is disobedient if it is possible that a_i was recommended but s_i should not be observed if every player were honest and obedient.

The incentive constraints $(*)$ incorporate all deviations by a player assuming others are honest and obedient. At disobedient histories, this assumption is contradicted by a player's signal. To deal with such histories, we make the following definitions.

Definition 4.3. A *sequential communication strategy* is a pair $(\sigma, \widetilde{\sigma})$ such that $\sigma \in \Delta(A \times A_0)$ is a correlated strategy and for every history (a_i, s_i), $\widetilde{\sigma}(\cdot|a_i, s_i)$ is a *positive* measure on $A_0 \times A$ such that given ζ,[15]

$$\sum_{(b,s_{-i})} \widetilde{\sigma}(\zeta, b|a_i, s_i) \Pr(s|b) \;=\; \sigma(a_i, \zeta) \;=\; \sum_{a_{-i}} \sigma(a, \zeta),$$

where $b \in A$.[16] Denote by Σ the set of all sequential communication strategies.

The condition above requires every player to trust that neither the organizer nor the monitoring techonology ever deviates.[17] We only ask of $\widetilde{\sigma}$ to be a positive measure because the product $\widetilde{\sigma} \Pr$ above is considered *proportional* to a player's beliefs. Specifically, player i believes that b was played, that ζ will be traded, and that players observed s with probability proportional to $\widetilde{\sigma}(\zeta, b|a_i, s_i) \Pr(s|b)$ after observing (a_i, s_i). Although $\widetilde{\sigma} \Pr$ does not add up to one on (ζ, s_{-i}, b), it is still interpretable as a player's conditional beliefs by a simple normalization, since the above condition ensures that $\widetilde{\sigma}$ adds up to some positive number on histories with $\sigma(a_i) > 0$. (These are the only histories of interest, as will be argued shortly, see Footnote 19.)

Notice that after a history (a_i, s_i), it is possible that a player ascribes positive probability not only to $b_i \neq a_i$, but also to more than one such b_i. This is interpreted as imperfect recall. Thus, we rely on S_i to capture all of a player's private information, including his memory.[18]

Definition 4.4. A sequential communication strategy $(\sigma, \widetilde{\sigma}) \in \Sigma$ is called a *sequential communication equilibrium* (SCE) if σ satisfies $(*)$ and $\widetilde{\sigma}$ satisfies

$$\sum_{(\zeta,b,s_{-i})} \widetilde{\sigma}(\zeta, b|a_i, s_i) \Pr(s|b)[v_i(b, \zeta(s_i, s_{-i})) - v_i(b, \zeta(\widehat{s}_i, s_{-i}))] \;\geq\; 0 \qquad (**)$$

for every disobedient history (a_i, s_i) and every $\widehat{s}_i \in S_i$.

[15] If the support of σ is not finite, then simply replace the index ζ with the Borel subsets of A_0.

[16] For every $j \in I$, the action b_j is player j's *actual* behavior after the recommendation a_j.

[17] The beliefs above are clearly limits of a sequence of Bayesian beliefs from correlated strategies with full support for each history. However, the set of all beliefs in a sequential communication strategy may not always arise from the same sequence. To justify, suppose for example that there are three players, only one of which has more than one action available. The other two players have perfectly correlated private signals. Requiring all beliefs to be obtained by the same sequence of full-support strategies implies that the two players hold identical beliefs after a zero probability event, which seems unreasonable unless there is only one possible conclusion from the common evidence.

[18] To obtain *perfect recall*, suppose that for every i, $S_i = A_i \times T_i$, and $\Pr(s|a) = P(t|a)$ (for some $P : A \to \Delta(T)$) if $s_i = (a_i, t_i)$ for every i, otherwise $\Pr(s|a)$ equals zero.

A sequential communication equilibrium involves a correlated strategy together with conditional beliefs at disobedient histories that make truthful reporting incentive compatible. By assumption, beliefs on the path of play follow Bayes' rule.

Condition ($**$) requires that for every disobedient history (a_i, s_i) there exist beliefs that make player i willing to report truthfully his privately observed signal. Since at such history player i *knows* that someone deviated, "all bets are off" as to the behavior of others, and player i is allowed to believe, for instance, that several players deviated (indeed, this might be the only hypothesis consistent with his evidence),[19] or that he himself deviated, as long as his beliefs are consistent with his evidence. A player does not believe that others will misreport their signals. This restriction on beliefs is without loss of generality by the *revelation principle*: for any mechanism that involves lying, there is another that is payoff-equivalent and involves truth-telling. Since ($*$) and ($**$) require players to always prefer honesty, this assumption is self-fulfilling.

Requiring σ to satisfy ($*$) implies that every individual, assuming that all other players are honest and obedient, prefers to obey the mediator as well as report truthfully. The requirement that others be honest when an individual contemplates deviating includes disobedient histories, i.e., histories where others realize that someone must have deviated. The incentive constraints ($*$) fail to impose the requirement that other players be honest on disobedient histories, yet they rely on such an assumption.

Condition ($**$) imposed on $\widetilde{\sigma}$ fills this void. On disobedient histories, players may hold any beliefs consistent with the organizer's strategy and the monitoring technology. Any other history off the path of play (i.e., (a_i, s_i) with $\sigma(a_i) = 0$) is *irrelevant*.[20]

We may now define the team's *metering problem*, explained below.

$$
\begin{aligned}
[\![\mathbf{v}|\Pr]\!](z) \;:=\; \sup_{(\sigma,\widetilde{\sigma})\in\Sigma} \sum_{(a,\zeta)} \sigma(a,\zeta)v(a,\zeta) \quad \text{s.t.} & \\
\sum_{(a_{-i},\zeta)} \sigma(a,\zeta)[v_i(a,\zeta) - v_i(b_i,a_{-i},\zeta|\rho_i)] &\geq 0 \\
\sum_{(\zeta,b,s_{-i})} \widetilde{\sigma}(\zeta,b|a_i,s_i)\Pr(s|b)[v_i(b,\zeta(s_i,s_{-i})) - v_i(b,\zeta(\widehat{s}_i,s_{-i}))] &\geq 0 \\
\sum_{(a,\zeta,s)} \sigma(a,\zeta)\Pr(s|a)\zeta(s) &= z.
\end{aligned}
$$

[19] Consider a history (a_i, s_i) such that s_i has positive probability only if two players deviate.

[20] If $\sigma(a_i) = 0$ then consistency between σ and $\widetilde{\sigma}$ requires that $\widetilde{\sigma}(a_i|a_i, s_i) = 0$ for any s_i, therefore the constraints in ($**$) are automatically satisfied.

The team's metering problem is a *linear program*. Its objective is to maximize the team's welfare by choosing a sequential communication strategy $(\sigma, \widetilde{\sigma}) \in \Sigma$.[21] The last constraint requires trades to add up to some given $z \in \mathbb{R}^\ell$ on average, as usual.

The first family of incentive constraints is indexed by team members i, individual action pairs $a_i, b_i \in A_i$, and reporting strategies $\rho_i : S_i \to S_i$. According to the constraints, obeying the mediator and reporting private signals truthfully must be incentive compatible for every player, assuming that everyone else is honest and obedient. Therefore, a player's utility from any possible deviation b_i and reporting strategy ρ_i cannot exceed the utility of playing a_i when recommended to do so by the mediator and reporting signals honestly. (This is just condition $(*)$.)

The second family, called *sequential reporting constraints*, is indexed by histories (a_i, s_i) and private signals $\widehat{s}_i$. For any history (a_i, s_i), a player's utility from reporting that s_i was privately observed must exceed his utility from any other possible report $\widehat{s}_i$, where expected utility is calculated with beliefs $\widetilde{\sigma}(\cdot|a_i, s_i)$ such that $(\sigma, \widetilde{\sigma}) \in \Sigma$.

The beliefs $\widetilde{\sigma}$ are chosen arbitrarily apart from their consistency with σ, and the sequential reporting constraints apply also on the path of play given σ. Previously, it was argued that Bayes' rule should determine players' conditional beliefs at any history (a_i, s_i) on the path of play, assuming players are honest and obedient. However, the sequential reporting constraints makes no such restriction. Fortunately, this is already taken into account by the constraints on σ, as the next result shows.

Lemma 4.5. *If $(\sigma, \widetilde{\sigma})$ is a feasible solution of the metering problem then so is $(\sigma, \widetilde{\sigma}')$, where $\widetilde{\sigma}' \Pr = \sigma$ on the path of play and $\widetilde{\sigma}' = \widetilde{\sigma}$ elsewhere.*

Proof. We must show that $\widetilde{\sigma}'$ satisfies the sequential reporting constraints at every history (a_i, s_i) such that $\Pr(a_i, s_i|\sigma) > 0$, where

$$\widetilde{\sigma}'(\zeta, b|a_i, s_i) \;=\; \begin{cases} \sigma(a_i, b_{-i}, \zeta) / \sum_{s_{-i}} \Pr(s|b) & \text{if } b_i = a_i \text{ and } \sum_{s_{-i}} \Pr(s|b) > 0, \\ 0 & \text{otherwise.} \end{cases}$$

For any player i and private signal $\widehat{s}_i \in S_i$, define $\widehat{\rho}_i(s_i) = \widehat{s}_i$ and $\widehat{\rho}_i(s_i') = s_i'$ for any other private signal s_i'. From the constraints on σ for $\widehat{\rho}_i$ at $b_i = a_i$, it follows that

$$\sum_{(a_{-i},\zeta)} \sigma(a, \zeta)[v_i(a, \zeta) - v_i(a, \zeta|\widehat{\rho}_i)] \;\geq\; 0,$$

[21]Several potentially binding constraints are imposed on $(\sigma, \widetilde{\sigma})$ by requiring it to belong to Σ. They are left implicit in this requirement for the sake of expositional clarity, and do not disrupt the linear nature of the metering problem.

but since $\widehat{\rho}_i$ involves truthful reporting except at s_i, this becomes

$$\sum_{(a_{-i},\zeta,s_{-i})} \sigma(a,\zeta)\Pr(s|a)[v_i(a,\zeta(s)) - v_i(a,\zeta(\widehat{s}_i,s_{-i})] \;\geq\; 0.$$

By definition, replacing σ with $\widetilde{\sigma}'$ on the path of play preserves this inequality. □

On disobedient histories, we allow for any consistent beliefs adapted to Pr. Although the sequential reporting constraints require honesty at all histories, including ones that are irrelevant (i.e., (a_i, s_i) such that $\sigma(a_i) = 0$), this does not impose any effective restrictions on a team's problem, since consistency requires that $\widetilde{\sigma}(a_i|a_i, s_i) = 0$, which automatically satisfies the sequential reporting constraints.

Lemma 4.6. *The set of sequential communication equilibria is closed and convex.*

Proof. Closedness follows because the incentive constraints impose weak inequalities and any convergent sequence of sequential communication equilibria must eventually have the same set of disobedient histories. As for convexity, take any two sequential communication equilibria, $(\sigma,\widetilde{\sigma})$ and $(\mu,\widetilde{\mu})$. For any $\alpha \in (0, 1)$, it is immediate that $\alpha\sigma + (1 - \alpha)\mu$ satisfies the first-stage incentive constraints, which impose obedience as well as honesty on the path of play. A history (a_i, s_i) is disobedient according to $\alpha\sigma+(1-\alpha)\mu$ if either $\sigma(a_i) > 0$ or $\mu(a_i) > 0$ and both $\Pr(a_i, s_i|\sigma) = \Pr(a_i, s_i|\mu) = 0$. If only $\sigma(a_i) > 0$ (or only $\mu(a_i) > 0$) then let conditional beliefs be given by $\widetilde{\sigma}(\cdot|a_i, s_i)$ (or $\widetilde{\mu}(\cdot|a_i, s_i)$), so the incentive constraint at (a_i, s_i) is satisfied. If both $\sigma(a_i) > 0$ and $\mu(a_i) > 0$ then any convex combination of $\widetilde{\sigma}(\cdot|a_i, s_i)$ and $\widetilde{\mu}(\cdot|a_i, s_i)$ satisfies the incentive constraint at (a_i, s_i), since it is satisfied by each individually. Therefore, all the incentive constraints at disobedient histories according to $\alpha\sigma + (1 - \alpha)\mu$ are satisfied with these beliefs, completing the proof. □

Proposition 4.7. *The team's utility function* $[\![\mathbf{v}|\Pr]\!]$ *is concave in* z*.*

This proposition, a direct consequence of Lemma 4.6, implies that we may naturally feed the team's utility function into a general equilibrium problem involving teams in the spirit of Rahman (2005b). In other words, a team's game may be subsumed in the background, and summarized by the representative team's utility function above.

Finally, the constraint set in the metering problem is nonempty. Indeed, by letting $\zeta : S \to \mathbb{R}^\ell$ be a constant function of $s \in S$, any correlated equilibrium is also a sequential communication equilibrium. Since correlated equilibria exist (see, for instance, Myerson, 1997), it follows that the set of feasible SCE is nonempty.

4.3 Choosing a Monitor

Looking back at Example 4.2, the reason for w being unattainable is simply that player 2's preferences fail to satisfy a "separation property." Let us amend player 2's preferences until a version of this property is satisfied.

Suppose first that player 2's preferences are given by

$$v_2(w,z) \;=\; 1-z \;=\; 1-v_2(s,z).$$

For w to be implementable, we still require that $\zeta(g) \geq \zeta(b) + \frac{1}{2}$ in order to satisfy player 1's incentive constraint. Player 2's constraints look like $1-\zeta(g) \geq 1-\zeta(b)$ and $\zeta(b) \geq \zeta(g)$, both of which coincide in their simplification as $\zeta(b) \geq \zeta(g)$. As long as this condition is satisfied, honesty by player 2 is incentive compatible. However, this is inconsistent with player 1's constraint, therefore w is not implementable.

If player 2's preferences are instead given by

$$v_2(w,z) \;=\; z \;=\; 1-v_2(s,z),$$

then w is implementable. Indeed, player 2's incentive constraints now require that $\zeta(g) \geq \zeta(b)$, a condition that is implied by player 1's incentive constraint. Therefore, player 2's honesty is not a binding constraint. This follows because player 2's preferences separate with player 1's effort in line with honest reporting in such a way that they do not conflict with player 1's incentive constraint.

Of course, if player 2 was some disinterested party, who happened to be indifferent amongst the team's provision of the public good, then he would be happy to report truthfully, thus facilitating the costless implementation of incentive contracts. This simple argument seems to be a generally useful rule of thumb when looking for a suitable monitor: all else equal, it doesn't hurt if he is a disinterested party. (However, this option may not always be available, or optimal—all else might not be equal.)

Another way to guarantee honesty on the part of the monitor is by introducing private contracts. The next example assumes that the team has access to the incentive good.

Example 4.8. The environment is almost identical to the previous one except that now there are two goods, one public (denoted z_0) and one private (the incentive good, denoted z_i). Preferences for player 1 over the public good are as usual, and for player 2 are given by

$$v_2(w,z_0) \;=\; 1-z_0 \;=\; 1-v_2(s,z_0).$$

Hence, without the incentive good, w is not implementable. Preferences over the incentive good are additively separable from those over the public good, and marginal utility over the incentive good equals one.

One might think that the way to induce player 2 to be honest, assuming that $\zeta_0(g) \geq \zeta_0(b) + \frac{1}{2}$ so that player 1 wants to play w, is by paying him with a private contract $(\zeta_2(g), \zeta_2(b))$. Player 2's incentive constraints then look like

$$\begin{aligned} 1 - \zeta_0(g) + \zeta_2(g) &\geq 1 - \zeta_0(b) + \zeta_2(b) \\ \zeta_0(b) + \zeta_2(b) &\geq \zeta_0(g) + \zeta_2(g), \end{aligned}$$

which imply that $\zeta_0(b) \geq \zeta_0(g)$, rendering w impossible to implement yet again. The alternative approach—paying player 1 in units of the incentive good—successfully implements w, however. Indeed, consider the following contractual arrangement: $\zeta_0(g) = \zeta_0(b)$ and $(\zeta_1(g), \zeta_1(b)) = (\frac{1}{2}, 0)$, with $\zeta_2 \equiv 0$. It is clear that player 2 has the incentive to be honest. Player 1 is now indifferent between w and s, so in particular he is willing to play w, as required.

Although private contracts allow for honesty to be incentive compatible, it is crucial that the way in which this takes place is by rewarding active players with the incentive good and making inactive observers indifferent over their reports with contracts denominated in the public good. If the observers have to incur monitoring effort, we might resort to loyalty-testing contracts as in Section 2.2.

This suggests two dimensions of quality in a monitor:

- the precision of the monitor's information relative to his monitoring effort;
- the alignment between the monitor's preferences and the team's tastes.

The examples illustrate a potential trade-off between these two dimensions of quality, as well as convey the message that Holmström's two questions quoted in Section 1 have interrelated answers. In essence, *the choice of monitors is largely determined by how output should be shared in order to provide all members of the team with the best incentives to perform.*

Example 4.2 presents a team where sequential reporting is a binding constraint. The linear programming formulation of Section 4.2 hints towards the use of duality for finding a team's shadow cost of sequential reporting.

Consider the following linear program, which is *dual* to the team's metering problem:

$$\inf_{\lambda,\widetilde{\lambda}\geq 0,\beta,\mu,\nu} \quad \nu + \mu z \quad \text{s.t.}$$

$$v(a,\zeta) + \sum_{(i,a_i,b_i,\rho_i)} \lambda_i(a_i,b_i,\rho_i)[v_i(a,\zeta) - v_i(b_i,a_{-i},\zeta|\rho_i)] -$$

$$\mu \sum_{s\in S} \Pr(s|a)\zeta(s) + \sum_{(i,s_i)} \beta_i(\zeta|a_i,s_i) \;\leq\; \nu$$

$$\sum_{(\widehat{s}_i,s_{-i})} \widetilde{\lambda}_i(\widehat{s}_i|a_i,s_i) \Pr(s|b)[v_i(b,\zeta(s_i,s_{-i})) - v_i(b,\zeta(\widehat{s}_i,s_{-i}))] \;\leq\; \beta_i(\zeta|a_i,s_i) \sum_{s_{-i}} \Pr(s|b).$$

The dual objective is to minimize an expenditure $\nu + \mu z$, as usual. The first family of constraints, indexed by (a,ζ), arises from interpreting σ as multipliers, whereas the second family, indexed by (a_i,s_i) and (ζ,b), arises from $\widetilde{\sigma}$.

The variable ν is associated with the probability constraint on σ, μ is associated with the resource constraint, β is associated with the constraints requiring consistency between $\widetilde{\sigma}$ and σ, $\widetilde{\lambda}$ is associated with the sequential reporting constraints, and finally λ is associated with condition $(*)$.

Lemma 4.9. $\beta \leq 0$ *at any dual solution.*

Proof. Let $\widetilde{\lambda} = 0$, so that $\beta = 0$ is feasible whatever the value of the other variables. Since any feasible $\beta \leq 0$ leads to a lower value of ν than $\beta = 0$, the result follows. □

Ignoring the sequential reporting constraints leads to *communication equilibrium*. The lemma shows just how adding such constraints over and above communication equilibrium is potentially binding. The next result provides a criterion for contracts to not be implementable in SCE.

Proposition 4.10. *A correlated strategy* $\sigma \in \Delta(A \times A_0)$ *is* not *part of a SCE if there is a player* i *and signals* $s_i, \widehat{s}_i \in S_i$ *such that for some* $a_i \in A_i$ *with* $\sigma(a_i) > 0$*,*

$$\sum_{(\zeta,s_{-i})} \sigma(a_i,\zeta) \Pr(s|b)[v_i(b,\zeta(s)) - v_i(b,\zeta(\widehat{s}_i,s_{-i}))] \;<\; 0$$

for every $b \in A$ *with* $\sum_{s_{-i}} \Pr(s|b) > 0$.

Proof. This is an immediate consequence of Farkas' Lemma (see Rockafellar, 1970, page 200, for example) applied to the sequential reporting constraints. □

For instance, in Example 4.2 any σ that would induce player 1 to play w fails to satisfy the condition above. From Proposition 4.10 we obtain the following *necessary* condition for an incentive contract ζ to be implementable in sequential communication equilibrium with a *pure* trading strategy: for every player i and any pair of signals $s_i, \widehat{s}_i \in S_i$, there exists $b \in A$ with $\sum_{s_{-i}} \Pr(s|b) > 0$ such that

$$\sum_{s_{-i}} \Pr(s|b)[v_i(b,\zeta(s)) - v_i(b,\zeta(\widehat{s}_i, s_{-i}))] \;\geq\; 0.$$

This is the "separation property" alluded to before.

4.4 Optimum Contracts via Duality

Next, it will be argued that the broad contractual structure derived in Section 2.2 generalizes to any game where private monitoring effort is worthwhile. Specifically, we will show that these so-called *loyalty testing contracts* solve the metering problem in the presence of the incentive good.

For completeness, we begin with an example that modifies the team's monitoring technology of Section 2.2 as follows:

$$\Pr(b|\text{shirk,shirk}) \;=\; \Pr(b|\text{shirk,work}) \;=\; 1/2$$
$$\Pr(g|\text{monitor,work}) \;=\; 3/4, \qquad \Pr(g|\text{monitor,shirk}) \;=\; 2/3.$$

If Robinson shirks then signal realizations are equally likely. If Robinson monitors and Friday works then the signal w realizes with higher probability than if Friday shirks. The key difference between this example and the one in Section 2.2 is that here every signal realizes with positive conditional probability.

We will find contracts that induce Friday to work with positive probability. Once again, we resort to contracts contingent on effort recommendations and loyalty tests.

With the notation of Section 2.2, Robinson's expected payment if he tells the truth and obediently monitors equals

$$\mu[2 + \tfrac{3}{4}\zeta_r(g) + \tfrac{1}{4}\zeta_r(b)] + (1-\mu)[-1 + \tfrac{2}{3}\xi_r(g) + \tfrac{1}{3}\xi_r(b)].$$

If Robinson is disobedient and dishonest by reporting the signal that suits him most, then his payoff will be

$$3\mu + \max\left\{\mu\zeta_r(g) + (1-\mu)\xi_r(g), \mu\zeta_r(b) + (1-\mu)\xi_r(b)\right\}.$$

This leads to the following two inequalities:

$$\begin{aligned} -1 + \mu\tfrac{1}{4}[\zeta_r(b) - \zeta_r(g)] &\geq (1-\mu)\tfrac{1}{3}[\xi_r(g) - \xi_r(b)] \\ -1 + \mu\tfrac{3}{4}[\zeta_r(g) - \zeta_r(b)] &\geq (1-\mu)\tfrac{2}{3}[\xi_r(b) - \xi_r(g)]. \end{aligned}$$

The system of equations above is solved for $0 < \mu < 1$ by $\zeta_r(b) = \xi_r(g) = 0$ and

$$\zeta_r(g) = 12/\mu, \qquad \xi_r(b) = 12/(1-\mu).$$

With this arrangement, Robinson finds it conditionally optimal to report truthfully upon receiving private information as well as choosing to monitor when it is recommended to him. Indeed, using Bayes' rule, it is easy to verify that the expected utility to Robinson from telling the truth given any signal realization is greater than the expected utility associated with any other reporting strategy. Once Robinson's incentives are aligned, it is a standard exercise to induce Friday to be willing to work.

It is not always possible to appropriately align a monitor's incentives, though. For instance, consider a team with only one active player who is asked to monitor his own effort, i.e., his payment depends on his own report. In this case, there is simply no way to induce the active player to report his effort truthfully unless it already is incentive compatible for him to work without contractual payments. Therefore, private monitoring is useful only if the monitor is monitoring other individuals, it is useless if he has to monitor himself.

Loyalty-testing describes a general approach for inducing monitors to exert effort as well as report truthfully whenever possible. By the previous paragraph, "public" convex independence does not suffice to attain a given correlated strategy. (Another example to this effect is the action profile (monitor,work) in Section 2.2, which is not attainable.)

To begin with, let us analyze *communication equilibria*, thereby ignoring sequential reporting constraints. With the insight of Section 3.4, we obtain the following primal:

$$[\![\mathbf{v}||\Pr]\!](z) := \sup_{\sigma\geq 0,\xi} \sum_{(i,a)} \sigma(a)v_i(a) + \sum_{(i,a,s)} \xi_i(a,s)\Pr(s|a) \quad \text{s.t.}$$

$$\sum_{a_{-i}} \sigma(a)[v_i(b_i,a_{-i}) - v_i(a)] + \sum_{(a_{-i},s)} \xi_i(a,\rho_i(s_i),s_{-i})\Pr(s|b_i,a_{-i}) - \xi_i(a,s)\Pr(s|a) \leq 0$$

$$\sum_{(i,a,s)} \xi_i(a,s)\Pr(s|a) = z$$

$$\sum_{a\in A} \sigma(a) = 1$$

The primal objective is to maximize welfare, as usual, by choosing a correlated strategy σ and probability-weighted payments ξ, subject to the following constraints. Firstly, obedience to the organizer's recommendations as well as truthful reporting must be incentive compatible on the path of play. Secondly, expected payments must add up to available resources, and finally, σ must add up to one.

The dual is given by the following linear program:

$$\begin{aligned} [\![\mathbf{v} \| \Pr]\!](z) &= \inf_{\lambda \geq 0, \mu, \nu} \quad \nu + \mu z \quad \text{s.t.} \\ \sum_{i=1}^{n} v_i(a) + \sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i)[v_i(a) - v_i(b_i, a_{-i})] &\leq \nu \\ \Pr(s|a) + \sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i)[\Pr(s|a) - \Pr(s|b_i, a_{-i}, \rho_i)] &= \mu \Pr(s|a), \end{aligned}$$

where given $a \in A$, $\rho_i : S_i \to S_i$, and $s \in S$,

$$\Pr(s|a, \rho_i) := \sum_{\widehat{s}_i \in \rho_i^{-1}(s_i)} \Pr(\widehat{s}_i, s_{-i}|a).$$

The dual minimizes expenditure by choosing λ (the multipliers on the primal incentive constraints), μ (the multiplier on the resource constraint), and ν (the multiplier on the probability constraint).

The value $\Pr(s|a, \rho_i)$ captures the conditional probability that s will be reported given that a is played, that player i chooses the reporting strategy ρ_i, and all other players choose to report truthfully. Since ρ_i^{-1} partitions S_i, it is immediate that $\Pr(s|a, \rho_i)$ is a well-defined conditional probability, with $\sum_s \Pr(s|a, \rho_i) = 1$ for every ρ_i.

With the same proof as for Lemma 3.8, we obtain the following result and definition.

Lemma 4.11. *If (λ, μ, ν) solves the dual then $\mu = 1$. Therefore,*

$$\sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|a) = \sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|b_i, a_{-i}, \rho_i)$$

for every $i \in I$, $a \in A$, and $s \in S$, at any dual solution.

Definition 4.12. A private monitoring technology Pr satisfies *convex independence* if for every player $i \in I$ and action $a_i \in A_i$,

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

where τ_i is the truthful reporting strategy and $\Pr[b_i, \rho_i] : A_{-i} \to \Delta(S)$ is the family of conditional probability vectors for s given (b_i, a_{-i}, ρ_i), indexed by a_{-i}.

If Pr satisfies convex independence then for every (i, a_i),

$$\sum_{(b_i \neq a_i, \rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|a) \;=\; \sum_{(b_i \neq a_i, \rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|b_i, a_{-i}, \rho_i),$$

and $\lambda \geq 0$ implies that $\lambda = 0$, except on $b_i = a_i$, where there always exist nonzero weights that satisfy the equality above. Nonetheless, at such λ the left-hand side of the first dual constraints equals zero, the key step in proving Theorem 3.10.

As with public monitoring, convex independence of a private monitoring technology suffices for any team to overcome its incentive constraints, as well as being the weakest sufficient condition. The next result formalizes this statement. Its proof is identical to that of Theorem 3.10, therefore omitted.

Theorem 4.13. *The private monitoring technology* Pr *satisfies convex independence if and only if*

$$[\![\mathbf{v}|| \Pr]\!](0) \;=\; \max\{v(a) : a \in A\}$$

for every $\mathbf{v} \in \mathbb{R}^{nA}$*, where* $v := \sum_i v_i$ *and* $n > 1$*.*

Intuitively, convex independence captures the idea that a player's action can be detected whether or not he chooses to misreport his private signal. By Theorem 4.13, convex independence is equivalent to the statement that there are contractual payments that relax a team's incentive constraints.

For instance, the private monitoring technology in Section 2.2 satisfies convex independence. Next, we will describe what it is about Pr that makes some profiles attainable but not others. Following Section 3.4, call a correlated strategy $\sigma \in \Delta(A)$ *approachable* if there is a feasible solution to the primal (σ, ξ) for some ξ, and call it *attainable* if there is a feasible solution to the primal (σ, ξ) such that $\sigma(a) = 0$ implies $\xi_i(a, s) = 0$ for every (i, s). Once again, the next result has the same proof as its public monitoring counterpart, Corollary 3.11.

Corollary 4.14. *If* Pr *satisfies convex independence then every correlated strategy is approachable.*

Definition 4.15. Pr satisfies *convex independence at* $\sigma \in \Delta(A)$ if for every $i \in I$ and $a_i \in A_i$ such that $\sigma(a_i) > 0$,

$$\sigma[a_i] \Pr[a_i, \tau_i] \notin \text{conv}\{\sigma[a_i] \Pr[b_i, \rho_i] : b_i \neq a_i\},$$

where $\sigma[a_i] \Pr[b_i, \rho_i](a_{-i}) := \sigma(a) \Pr(b_i, a_{-i}, \rho_i)$ and $\Pr(b_i, a_{-i}, \rho_i) \in \Delta(S)$.

Lemma 4.16. Pr *satisfies convex independence if and only if* Pr *satisfies convex independence at every* σ *with full support.*

Proof. Pr satisfies convex independence if and only if given (i, a_i) for which there exists $\lambda_i(a_i) > 0$ such that

$$\sum_{(b_i \neq a_i, \rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|a) = \sum_{(b_i \neq a_i, \rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|b_i, a_{-i}, \rho_i)$$

for every (a_{-i}, s), there exists $\widehat{a}_i$ such that

$$\sum_{(b_i \neq \widehat{a}_i, \rho_i)} \lambda_i(\widehat{a}_i, b_i, \rho_i) \Pr(s|\widehat{a}_i, a_{-i}) \neq \sum_{(b_i \neq \widehat{a}_i, \rho_i)} \lambda_i(\widehat{a}_i, b_i, \rho_i) \Pr(s|b_i, a_{-i}, \rho_i),$$

for some (a_{-i}, s). But this condition is independent of whether or not both sides of each equation are multiplied by $\sigma(a)$, $\sigma(\widehat{a}_i, a_{-i})$, respectively, if both are positive. □

Yet again, the next result has the same proof as its public version, Proposition 3.13.

Proposition 4.17. *If* Pr *satisfies convex independence at* σ *then* σ *is attainable.*

In Section 2.2, Pr fails to satisfy convex independence at (monitor,work), and as such is unattainable. (It is approachable, since Pr satisfies convex independence.)

As with public monitoring, for any approachable correlated strategy there is a sequence of correlated strategies converging to it and bounded contractual payments that define communication equilibria.[22]

If Pr satisfies convex independence then by Proposition 4.17 this sequence of correlated strategies may be taken to have full support. Therefore, by Lemma 4.5 and since every history is on the path of play (except perhaps histories that are only consistent with a player disobeying his own recommendation, which are already accounted for in $(*)$), it follows that this sequence of correlated strategies together with their respective implementing contracts define not just a communication equilibrium, but also a sequential communication equilibrium. Therefore, any correlated strategy with full support is *sequentially attainable*, i.e., there are contracts that make it a SCE. Moreover, we have the following result.

Call a correlated strategy σ *sequentially approachable* if there is a sequence of sequentially attainable correlated strategies $\{\sigma_m\}$ such that $\sigma_m \to \sigma$ as $m \to \infty$.

[22]Lemma 3.14 and Proposition 3.15 also apply to this primal with identical proofs, as well as the contractual structure discussed at the end of Section 3.4.

Proposition 4.18. *If* Pr *satisfies convex independence then every correlated strategy is sequentially approachable.*

We conclude this section by defining the metering problem with a monitoring technology that does not necessarily satisfy convex independence, for which sequential reporting constraints must be explicitly introduced. With the insight from Section 3.4, we will formulate a linear program and use duality to understand it better. First, consider the following (nonlinear) program:

$$\llbracket \mathbf{v} | \Pr \rrbracket (z) \;:=\; \sup_{\sigma, \widetilde{\sigma} \geq 0, \zeta} \sum_{(i,a,s)} \sigma(a)[v_i(a) + \zeta_i(a,s)\Pr(s|a)] \quad \text{s.t.}$$

$$\begin{aligned}
\sum_{(a_{-i},s)} \sigma(a)[(v_i(b_i,a_{-i}) - v_i(a) + \zeta_i(a,\rho_i(s_i),s_{-i})\Pr(s|b_i,a_{-i}) - \zeta_i(a,s)\Pr(s|a)] &\leq 0 \\
\sum_{(a_{-i},b,s_{-i})} \widetilde{\sigma}(a,b|a_i,s_i)\Pr(s|b)[\zeta_i(a,\widehat{s}_i,s_{-i}) - \zeta_i(a,s)] &\leq 0 \\
\sum_{(b,s_{-i})} \widetilde{\sigma}(a,b|a_i,s_i)\Pr(s|b) - \sigma(a) &= 0 \\
\sum_{(i,a,s)} \sigma(a)\zeta_i(a,s)\Pr(s|a) &= z \\
\sum_{a \in A} \sigma(a) &= 1.
\end{aligned}$$

As usual, this problem's objective is to maximize the team's welfare by choosing a correlated strategy σ together with beliefs $\widetilde{\sigma}$ and incentive contracts $\zeta : I \times S \times A \to \mathbb{R}$, since, by linearity of utility from the incentive good, there is no use for randomization above and beyond choosing potentially different schemes for each $a \in A$.

The first family of constraints is a version of $(*)$ applied to this particular environment. The second family corresponds to $(**)$, where the summation is with respect to a_{-i} instead of ζ since now this indexes incentive contracts chosen by the organizer. The third family of constraints, indexed by (a, s_i), ensures that σ and $\widetilde{\sigma}$ constitute a sequential communication strategy. The penultimate constraint is the usual restriction on contractual resources, and finally the last constraint ensures that σ is a probability measure.

For the same reasons as those discussed in Section 3.4, the program above does not necessarily have a solution. For instance, this optimization applied to the example of Section 2.2 has no solution even though one can feasibly get arbitrarily close to any correlated strategy. (See Claim 3.7, which also applies to private monitoring.)

Consider the following linear program, called the *private primal*:

$$[\![\mathbf{v}|||\Pr]\!](z) \;:=\; \sup_{\sigma\geq 0,\xi,\widetilde{\xi}} \sum_{(i,a)}\sigma(a)v_i(a) + \sum_{(i,a,s)}\xi_i(a,s)\Pr(s|a) \quad \text{s.t.}$$

$$\sum_{a_{-i}}\sigma(a)[v_i(b_i,a_{-i}) - v_i(a)] + \sum_{(a_{-i},s)}\xi_i(a,\rho_i(s_i),s_{-i})\Pr(s|b_i,a_{-i}) - \xi_i(a,s)\Pr(s|a) \leq 0$$

$$\sum_{(a_{-i},b,s_{-i})}\Pr(s|b)[\widetilde{\xi}_i(a,b,\widehat{s}_i,s_{-i}|a_i,s_i) - \widetilde{\xi}_i(a,b,s|a_i,s_i)] \leq 0$$

$$\sum_{(b,s_{-i})}\widetilde{\xi}_i(a,b,\widehat{s}|a_i,s_i)\Pr(s|b) - \xi_i(a,\widehat{s}) = 0$$

$$\sum_{(i,a,s)}\xi_i(a,s)\Pr(s|a) = z$$

$$\sum_{a\in A}\sigma(a) = 1$$

The private primal is just like the problem in the previous page except for three details. Firstly, any occurrence of $\sigma(a)\zeta_i(a,s)$ has been replaced with $\xi_i(a,s)$ as in Section 3.4. Secondly, any occurrence of $\widetilde{\sigma}(a,b|a_i,s_i)\zeta_i(a,\widehat{s})$ has been replaced with $\widetilde{\xi}_i(a,b,\widehat{s}|a_i,s_i)$. Finally, the consistency condition between σ and $\widetilde{\sigma}$ (the latter no longer chosen) has been replaced with one between ξ and $\widetilde{\xi}$, indexed by $(a,s_i,\widehat{s})$.

The *private dual*, i.e., the dual of the private primal, is given by:

$$[\![\mathbf{v}|||\Pr]\!](z) \;=\; \inf_{\lambda,\widetilde{\lambda}\geq 0,\beta,\mu,\nu} \nu + \mu z \quad \text{s.t.}$$

$$\sum_{i=1}^{n} v_i(a) + \sum_{(a_i,b_i,\rho_i)}\lambda_i(a_i,b_i,\rho_i)[v_i(a) - v_i(b_i,a_{-i})] \;\leq\; \nu$$

$$\sum_{(a_i,b_i,\rho_i)}\lambda_i(a_i,b_i,\rho_i)[\Pr(s|a) - \Pr(s|b_i,a_{-i},\rho_i)] + \sum_{\widehat{s}_i}\beta_i(a,s|a_i,\widehat{s}_i) \;=\; (\mu-1)\Pr(s|a)$$

$$(\widehat{s}_i = s_i) \qquad \sum_{s_{-i}}\Pr(s|b)\beta_i(a,\widehat{s}|a_i,s_i) + \sum_{s_i'\neq s_i}\widetilde{\lambda}_i(s_i'|a_i,s_i)\Pr(s_i',\widehat{s}_{-i}|b) \;=\; 0$$

$$(\widehat{s}_i \neq s_i) \qquad \sum_{s_{-i}}\Pr(s|b)\beta_i(a,\widehat{s}|a_i,s_i) - \widetilde{\lambda}_i(\widehat{s}_i|a_i,s_i)\Pr(s_i,\widehat{s}_{-i}|b) \;=\; 0$$

The private dual minimizes expenditure as usual by choosing λ (the multipliers for $(*)$), $\widetilde{\lambda}$ (the multipliers for $(**)$), β (the multipliers for the consistency constraints between ξ and $\widetilde{\xi}$), μ (the multiplier on the resource constraint, and ν (the multiplier on the probability constraint).

The first family of constraints, indexed by $a \in A$, is identical to those in the previous dual that ignored sequential reporting constraints to find communication equilibria.

The second family, indexed by (i, a, s), would coincide with the previous dual if it weren't for the term $\sum_{\widehat{s}_i} \beta_i(a, s|a_i, s_i)$. If this term were equal to zero, then the private dual would be equivalent to the previous dual. The third family of constraints is indexed by $(a, b, \widehat{s})$ and histories (a_i, s_i) such that $\widehat{s}_i = s_i$, and the fourth family is indexed by $(a, b, \widehat{s})$ and histories (a_i, s_i) such that $\widehat{s}_i \neq s_i$.

Adding together the third and fourth families of constraints leads to the next result.

Lemma 4.19. *If $(\lambda, \widetilde{\lambda}, \beta, \mu, \nu)$ solves the dual then $\mu = 1$ and*

$$\sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|a) \;=\; \sum_{(b_i,\rho_i)} \lambda_i(a_i, b_i, \rho_i) \Pr(s|b_i, a_{-i}, \rho_i)$$

for every $i \in I$, $a \in A$, and $s \in S$.

Proof. Fix a player i, a pair of action profiles $a, b \in A$, a signal profile $\widehat{s}_{-i} \in S_{-i}$, and a history (a_i, s_i). Adding the fourth constraints across $\widehat{s}_i \neq s_i$, and then adding to this subtotal the associated third constraint, it follows that

$$\sum_{\widehat{s}_i \in S_i} \Pr(s_i|b) \beta_i(a, \widehat{s}|a_i, s_i) \;=\; 0,$$

where $\Pr(s_i|b) := \sum_{s_{-i}} \Pr(s|b)$. Since this must be true for any $b \in A$, and given s_i there exists b such that $\Pr(s_i|b) > 0$, it follows that $\sum_{\widehat{s}_i} \beta_i(a, \widehat{s}|a_i, s_i) = 0$. Adding the second family of dual constraints with respect to $s \in S$, we obtain

$$\sum_{s \in S} \sum_{\widehat{s}_i} \beta_i(a, s|a_i, \widehat{s}_i) \;=\; \sum_{s_i \in S_i} \sum_{(\widehat{s}_i, s_{-i})} \beta_i(a, \widehat{s}_i, s_{-i}|a_i, s_i) \;=\; 0,$$

where the second equality follows by the previous argument. The rest of the proof now mimics that of Lemma 3.8. □

With an identical proof to that of Theorem 3.10, we obtain the following result.

Theorem 4.20. *The private monitoring technology* Pr *satisfies convex independence if and only if*

$$[\![\mathbf{v}|||\Pr]\!](0) \;=\; \max\{v(a) : a \in A\}$$

for every $\mathbf{v} \in \mathbb{R}^{nA}$, where $v := \sum_i v_i$ and $n > 1$.

This statement differs from Theorem 4.13 in that here the value function on the left-hand side includes sequential reporting constraints. To some extent, this result is already incorporated in Theorem 4.13 since clearly $[\![\mathbf{v}|||\Pr]\!] \leq [\![\mathbf{v}||\Pr]\!]$ (one problem is more constrained than the other).

5 Conclusions

This paper has developed a formal approach to obtain optimum contracts in general environments via duality. Necessary and sufficient conditions on a public and private monitoring technology were derived for any team to overcome its incentive constraints and implement its first best outcome. Recommendation-contingent contracts were shown to often be useful (with public monitoring, they can attain outcomes with a smaller average budget) and sometimes crucial (with private monitoring, some outcomes might not otherwise be implementable) to a team's organization.

A shortcoming of loyalty-testing contracts is that they are not robust to collusion. For instance, if Robinson and Friday (in Section 2.2) could communicate with each other after receiving their recommendations and before taking actions, they would prefer to do so, allowing them both to shirk and receive payment. However, if Robinson experienced a positive cost of communicating this way (however small), the mediating principal could dissuade Robinson from doing so by having him monitor more Fridays, thereby "saturating away" Robinson's collusive incentives.

Although we assumed throughout that teams are subject only to expected resource constraints, as opposed to "ex post" budget balance, incorporating such constraints into the linear programming formulation developed here would not be difficult, and its effect on a team's problem would be comparable to imposing contractual boundedness as developed at the end of Section 3.4.

Using duality, convex independence was shown to characterize attainability of efficient outcomes. This is a popular condition in the literature on repeated games. Some of the most prominent contributions that use versions of this to establish folk theorems are Compte (1998), Kandori and Matsushima (1998), and recently Obara (2005).

Fudenberg et al. (1994) established a fruitful link between implementable outcomes in one-shot games with linear transfers and implementable outcomes in repeated games, namely that for sufficiently patient players, reward and punishment regimes in dynamic strategic interactions are interpretable as linear transfers in the "steady-state" stage game. The present paper might be thought of as forging the link of Fudenberg et al. (1994) by characterizing stage games (via convex independence) whose repeated version would lead to a folk theorem with public or private monitoring. Formalizing this link with the use of linear programming duality techniques as developed here will be the subject of further research by the author.

References

Alchian, A. and H. Demsetz (1972): "Production, Information Costs, and Economic Organization," *American Economic Review*, 62, 777–795.

Aumann, R. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

Compte, O. (1998): "Communication in Repeated Games with Imperfect Private Monitoring," *Econometrica*, 66, 597–626.

Fudenberg, D., D. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62, 997–1039.

Holmström, B. (1982): "Moral Hazard in Teams," *Bell Journal of Economics*, 13, 324–340.

Kandori, M. and H. Matsushima (1998): "Private Observation, Communication, and Collusion," *Econometrica*, 66, 627–652.

Makowski, L. and J. Ostroy (2003): "Competitive Contractual Pricing with Transparent Teams," mimeo.

Myerson, R. (1986): "Multistage Games with Communication," *Econometrica*, 54, 323–358.

——— (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. (2005): "Folk Theorem with Communication," mimeo.

Rahman, D. (2005a): "Competition with Local Mechanisms: Allocating Residual Claims and Control Rights," mimeo.

——— (2005b): "Contractual Pricing with Incentive Constraints," mimeo.

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.