

Incentive compatible contractible information[★]

Philip Bond

Kellogg School of Management, Northwestern University, Evanston IL 60208-2001, USA
(e-mail: p-bond@northwestern.edu)

Received: August 7, 2000; revised version: December 21, 2001

Summary. The paper seeks to characterize what information is always available for contracting, independent of the form of the contract and the probabilities of different states of nature. The paper denotes such information as *contractible*. It is established that it is possible to speak uniquely of *maximal contractible information*. Several characterizations are exhibited. In particular, it is shown that if either (a) punishments are bounded everywhere, or (b) deviations from truth-telling are either always or never detected, then maximum contractible information coincides with $\bigwedge_i \bigvee_{j \neq i} \mathcal{E}_j$ where \mathcal{E}_j is the information partition of agent j . An argument is given for why (b) may be expected to hold.

Keywords and Phrases: Contractible information, Incentive compatibility, Information partition, Common knowledge, Cross-reporting.

JEL Classification Numbers: D78, D82.

1 Introduction

Incentive constraints play a central role in understanding the restrictions imposed by private information (i.e. information uniquely possessed by single agents). Dual to this, much of the contract theory literature has assumed that information possessed by at least two agents can be freely contracted upon, regardless of the nature of the contract and the probability of different states.¹ For instance, in principal-agent models² effort is assumed private to the agent, but output is assumed to be observed by both the principal and agent and hence to be contractible.

* I thank Michael Chwe, Douglas Diamond, Lars Stole, Robert Townsend, Nicholas Yannelis and an anonymous referee for helpful comments.

¹ The exceptions are the implementation theory and incomplete contracting literatures.

² See, for instance, Mirrlees (1999).

This paper seeks to extend our understanding of what information is contractible in three ways. First, what information can be said to be contractible in more complex informational settings? For example, suppose that in a principal-agent setting the agent either has malaria, the flu, a cold, or is entirely healthy, but is unable to distinguish between flu and a cold. Suppose (for the sake of example) that malaria and the flu produce identical symptoms, and that the principal (but not the agent, who in this example lacks medical training) is able to identify these symptoms. So the principal knows when the agent is suffering from malaria or the flu, but not which, and knows when the agent either has a cold or is healthy, but not which. Is any aspect of the agent's health always contractible here? The answer turns out to depend on whether or not the size of punishments available is bounded.

Second, when there are more than two agents, is information private unless it is entirely shared by two agents? For instance, suppose effort has two components – duration and intensity. If the principal observes hours worked, and a coworker intensity, is the agent's effort contractible?

Third, is it possible to speak uniquely of maximal *contractible information*, in the sense of a uniquely defined maximal amount of information that is freely contractible? If so, how is such information characterized?

Before proceeding to the analysis, it is worth pausing to note how this paper differs from the work initiated by Crémer and McLean (1985)³ establishing circumstances in which incentive compatibility places no constraint on the ability of an uninformed outsider to extract surplus. This body of work is concerned with incentive compatibility's effect on what welfare levels are attainable, and depends on the assumptions of risk neutral agents and correlated information. In contrast, the present paper considers what *contracts* are unaffected by incentive concerns, and makes only weak assumptions about preferences, and none about stochastic structure.

I conclude the introduction with a note on methodology. This paper is a *mechanism design* paper, in the sense of being concerned with whether outcomes can be supported as the equilibria of some pre-specified game. It neglects the concerns of the parallel *implementation theory* literature about whether there also exist other (undesirable) equilibria of these games. It is well known (see Maskin, 1999) that the multiple equilibrium problem is important in complete information settings (i.e. those in which agents share all information) when the equilibrium concept employed is Nash. However, papers by – among others – Matsushima (1993), Arya et al. (1995), and Duggan (1997) suggest that the multiple equilibrium problem is of limited importance when information is incomplete. Moreover, in complete information settings Moore and Repullo (1988) find that multiple equilibria concerns can be avoided if one looks instead at implementation in subgame perfect equilibria. In this paper I appeal to this broad class of results for justification in considering the mechanism design problem in isolation.

The paper proceeds as follows. Section 2 gives a general specification of the problem to be analyzed. Section 3 establishes that maximal contractible information is uniquely defined. Section 4 presents several examples. Section 5 shows that

³ See also Crémer and McLean (1988), together with the related work of McAfee and Reny (1992).

if punishments are unlimited, maximum contractible information has an easy characterization. Section 6 establishes that the same characterization holds with limited punishments if information satisfies a certain restriction, and gives an argument for why that restriction is reasonable. Section 7 discusses the case in which preferences as well as probabilities are allowed to vary. Section 8 gives an alternative characterization of maximum contractible information. Section 9 relates the results of this paper to previous research.

2 Preliminaries

An economy is a quintuple $(\Omega, N, A, \{u_i\}, \{\mathcal{E}_i\})$ where:

- Ω is the state space, with $\omega \in \Omega$ a typical member. Let M be the (finite) cardinality of Ω .
- N is the (finite) set of agents.
- A is the (possibly infinite) set of outcomes.
- $u_i : A \times \Omega \rightarrow \Re$ is the utility mapping of agent $i \in N$, giving the utility of each agent given an outcome and state of the world. Agents are assumed to maximize expected utility.
- \mathcal{E}_i is a partition of Ω giving the information⁴ possessed by agent $i \in N$.

Also, let P be the set of probability mappings $p : \Omega \rightarrow (0, 1)$, where $p(\omega)$ is the probability of state ω . Thus $P \subset (0, 1)^M$.

The economy $(\Omega, N, A, \{u_i\}, \{\mathcal{E}_i\})$ and probability mapping $p \in P$ are assumed to be common knowledge to the agents N . We will characterize what information is always available for contracting, regardless of the probability mapping p . That is, if an outside observer sees the preferences and information of each agent in the economy, what information can that observer infer is *always* available for contracting, no matter what the probabilities are? In Section 7 we will also briefly consider the related question of what information is contractible independent of agents' preferences.

Throughout the paper, agents' preferences will be assumed to satisfy the following two properties:

Condition SP (Strict Preferences): *For all states $\omega \in \Omega$ and agents $i \in N$, there exist outcomes a_1 and a_2 such that $u_i(a_2; \omega) > u_i(a_1; \omega)$.*

Condition FD (Free Disposal): *For any finite subset $A_k = \{a_1, \dots, a_k\}$ of A , there exists an outcome $a_0 \in A$ such that $u_i(a_0, \omega) \leq u_i(a, \omega)$ for all agents $i \in N$, states $\omega \in \Omega$, and outcomes $a \in A_k$.*

Condition SP ensures that agents always have *strict preferences* over at least two outcomes, and serves to rule out degenerate cases in which information is contractible simply because agents face no meaningful choices. Condition FD represents a generalization of *free disposal of goods* in the following sense. Suppose

⁴ See, e.g., Osborne and Rubinstein (1994) for a description of how information can be represented by partitions of the state space.

that outcomes relate to the provision of (public or private) goods. Then provided all agents weakly prefer more goods to less in all states, the outcome a_0 just relates to taking the minimum provision of each good from the outcomes in A_k . Condition FD ensures that a central planner does not need to have too much information in order to effectively punish agents. Note that the condition says nothing about the *size* of punishments available, a point that will be returned to below.

Throughout the paper, \mathcal{E}_S will be used to denote the collective information $\bigvee_{i \in S} \mathcal{E}_i$ of a coalition $S \subset N$, and \mathcal{E}_{-i} will denote the collective information of all agents other than agent i , $\mathcal{E}_{N \setminus \{i\}}$. Also, $E_i(\omega)$ (respectively $E_S(\omega), E_{-i}(\omega)$) will denote the element of the partition \mathcal{E}_i (respectively $\mathcal{E}_S, \mathcal{E}_{-i}$) that contains the state ω .

Given a particular probability map p , incentive compatibility is defined as normal:

Definition 1 *A mapping $f : \Omega \rightarrow A$ is said to be p -incentive compatible⁵ if when probabilities are given by $p \in P$, there exists a mechanism $(\{\mathcal{M}_i\}, F : \times_i \mathcal{M}_i \rightarrow A)$ such that for all $\omega \in \Omega$ there is a Bayesian equilibrium $m(\omega)$ of the mechanism with $F(m(\omega)) = f(\omega)$.*

We are interested in what information can always be contracted upon, independent both of the desired outcomes and the particular probabilities of different states. Such information will be said to be *contractible*. Informally, information is contractible if an outside observer who is ignorant of the stochastic structure of the economy can nonetheless infer that the information is available for contracting. Formally,

Definition 2 *A partition \mathcal{G} of Ω is said to be contractible information if whenever $f : \Omega \rightarrow A$ is \mathcal{G} -measurable, then f is p -incentive compatible for all $p \in P$.*

Clearly it is never possible to contract on information not possessed by *any* agent. That is,

Lemma 1 *A partition \mathcal{G} is contractible information only if $\mathcal{G} \preceq \mathcal{E}_N$.*

Proof. Suppose not. Then for any $p \in P$ there exists a mapping $f : \Omega \rightarrow A$ and states $\omega_1, \omega_2 \in \Omega$ such that f is p -incentive compatible, $f(\omega_1) \neq f(\omega_2)$ and $E_N(\omega_1) = E_N(\omega_2)$. The latter implies that $E_i(\omega_1) = E_i(\omega_2)$ for all $i \in N$. But the revelation principle implies that there exists a direct mechanism – so

$$\begin{aligned} F(E_1(\omega_1), \dots, E_N(\omega_1)) &= f(\omega_1) \\ F(E_1(\omega_2), \dots, E_N(\omega_2)) &= f(\omega_2) \end{aligned}$$

for some function $F : \times_i \mathcal{E}_i \rightarrow A$. But since $f(\omega_1) \neq f(\omega_2)$ this gives a contradiction, completing the proof.

⁵ Note that incentive compatibility is sometimes referred to elsewhere as weak implementability.

3 Maximal contractible information

3.1 Punishments

Consider the example economies of Figures 1 and 2. The connected points depict the information partitions – so in Figure 1, $\mathcal{E}_1 = \{\{\omega_1\}, \{\omega_2, \omega_3\}, \{\omega_4\}\}$ and $\mathcal{E}_2 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$. The economy of Figure 1 corresponds to the health example given in the introduction. That is, the states $\omega_1, \dots, \omega_4$ correspond respectively to “malaria”, “flu”, “cold” and “healthy”. Agent 1 is the worker, and cannot distinguish between the flu and a cold. Agent 2 is the employer, who can spot the identical symptoms produced by malaria and the flu.




Figure 1. Information about health

Is the health of the worker available for contracting, independent of the form of the contract? The answer depends on the size of punishments available. To see this, consider the situation faced by the employer when she observes the symptoms associated with malaria and the flu (i.e. when her information is $\{\omega_1, \omega_2\}$). If she reports that the worker has no symptoms, the worker will know she is lying if he in fact has malaria (state ω_1), but not if he has the flu or a cold (states ω_2, ω_3), since in this case the worker is aware his illness might only be a cold. If it is possible to heavily punish the employer, then the threat of detection will be enough to keep her honest, and she will always tell the truth. In this case, contracting upon at least some aspects of the worker’s health (here, symptoms/no symptoms) is always possible. But if punishments are more limited, she may (depending on what is at stake) take the risk of detection and misreport the worker’s medical condition – so arbitrary contracts cannot be written on any aspect of the worker’s health.

The size of punishments is important here because of the availability of misreports that are detected sometimes, but not always. No such reporting deviations are available in the economy of Figure 2, and so – as we will see – the size of punishment available does not play a critical role.

With the above discussion in mind, at each state $\omega \in \Omega$ the agents N are partitioned into two groups $N_\infty(\omega)$ and $N_0(\omega)$ as follows: Agent i is in $N_\infty(\omega)$ if it holds that for all $x \in \mathfrak{R}$ there exists $a \in A$ such that $u_i(a, \omega) < x$. Otherwise




Figure 2. Alternative information about health

agent i is in $N_0(\omega)$. That is, $N_\infty(\omega)$ consists of those agents who at state ω can be punished *arbitrarily* harshly. For those agents in $N_0(\omega)$ for whom punishments are limited, define the maximum level of punishment $z_i(\omega)$ by

$$z_i(\omega) = \inf_{a \in A} u_i(a, \omega)$$

3.2 A candidate partition

Based on the above considerations about types of deviation and size of punishment available, I start by constructing a candidate partition for the maximum contractible information.

The revelation principle tells us that we can think in terms of agents truthfully reporting their information (i.e. elements of \mathcal{E}_i), and focuses attention on what incentives agents must be given to do so. Denote the report of each agent by $\tilde{E}_i \in \mathcal{E}_i$. If an agent's report is *always* inconsistent with the truthful reports of others (i.e. if $\tilde{E}_i \cap E_{-i}(\omega) = \emptyset$ for all $\omega \in E_i$, where E_i is agent i 's actual information), then it will be easy to deter untruthful reports by punishing the agent at these report combinations. Note that the punishments will not have to very large, since the deviating agent knows with certainty that he will be detected.

Report deviations that are only sometimes undetected will be harder to deter, and indeed may be impossible to do so for arbitrary outcome functions $f : \Omega \rightarrow A$. With this in mind, for each $i \in N$ define a binary relation \rightarrow_i on Ω by:

$$\omega_1 \rightarrow_i \omega_2 \text{ if and only if } \omega_2 \in E_{-i}(\omega_1)$$

So the relations \rightarrow_i define states which may be contractibly indistinguishable, since agent i can report $E_i(\omega_2)$ when her information is $E_i(\omega_1)$, and escape detection at least sometimes.

Whether they are or not depends, as we saw above, on the level of punishments available. With heavy punishments, occasional detection is a powerful deterrent – but with limited punishments this is not the case. Define a further set of binary

relations $\omega_1 \rightarrow_i \omega_2$ by

$$\omega_1 \rightarrow_i \omega_2 \text{ if and only if } \omega_1 \rightarrow_i \omega_2 \text{ and} \\ i \in N_0(\omega) \text{ whenever } \omega \in E_i(\omega_1) \text{ and } E_i(\omega_2) \cap E_{-i}(\omega) = \emptyset$$

So if $\omega_1 \rightarrow_i \omega_2$, agent i can report $E_i(\omega_2)$ at $E_i(\omega_1)$ knowing that this deviation will not always be detected, and in those states where it is punishments are limited.

Finally define the graph (Ω, \rightarrow) by

$$\omega_1 \rightarrow \omega_2 \text{ if and only if } \omega_1 \rightarrow_i \omega_2 \text{ for some } i \in N$$

Define \mathcal{E}^* as the partition induced by the components of the graph (Ω, \rightarrow) . Below it is shown that \mathcal{E}^* is the unique maximal contractible information partition.

First, note that the partition \mathcal{E}^* is coarser than the collective knowledge of all agents, \mathcal{E}_N .

Lemma 2 $\mathcal{E}^* \preceq \mathcal{E}_N$.

Proof. Take $\omega_1, \omega_2 \in \Omega$ such that $E_N(\omega_1) = E_N(\omega_2)$. Take any $i \in N$. Then $E_{-i}(\omega_1) = E_{-i}(\omega_2)$ and $E_i(\omega_1) = E_i(\omega_2)$. So certainly $\omega_1 \rightarrow_i \omega_2$. Moreover, since $E_i(\omega_1) \cap E_{-i}(\omega) \neq \emptyset$ for all $\omega \in E_i(\omega_1)$, it must also hold that $E_i(\omega_2) \cap E_{-i}(\omega) \neq \emptyset$ for all $\omega \in E_i(\omega_1)$. So $\omega_1 \rightarrow_i \omega_2$, completing the proof.

Moreover, \mathcal{E}^* is indeed contractible information:

Lemma 3 \mathcal{E}^* is contractible information.

Proof. Take an arbitrary \mathcal{E}^* -measurable function $f : \Omega \rightarrow A$ and probability mapping $p \in P$. We must show that f is p -incentive compatible. The proof is by construction, and consists of exhibiting a direct mechanism $F : \times_{i \in N} \mathcal{E}_i \rightarrow A$ such that

$$F((E_i(\omega))_{i \in N}) = f(\omega) \tag{1}$$

and such that for all $i \in N$ and $E_i, \tilde{E}_i \in \mathcal{E}_i$ the incentive constraint

$$\sum_{\omega \in E_i} p(\omega) u_i(f(\omega), \omega) \geq \sum_{\omega \in E_i} p(\omega) u_i\left(F\left(\tilde{E}_i, E_{-i}(\omega)\right), \omega\right) \tag{2}$$

holds.

First, choose $a_f \in A$ such that for $i \in N$ and $\omega, \omega' \in \Omega$, $u_i(a_f, \omega) \leq u_i(f(\omega'), \omega)$. Such a choice is always possible by Condition FD.

Next, construct a set of functions $F^i : (\times_{j \in N} \mathcal{E}_j) \times \mathcal{E}_i \rightarrow A$ as follows. First, if $\left(\left(\tilde{E}_j\right)_{j \in N}, E_i\right) \in (\times_{j \in N} \mathcal{E}_j) \times \mathcal{E}_i$ is such that $\bigcap_{j \in N} \tilde{E}_j \neq \emptyset$, then set $F^i\left(\left(\tilde{E}_j\right)_{j \in N}, E_i\right) = f\left(\bigcap_{j \in N} \tilde{E}_j\right)$. Note that this is a well-defined choice, since by assumption f is \mathcal{E}^* -measurable and hence \mathcal{E}_N -measurable by Lemma 2.

Over the remainder of $(\times_{j \in N} \mathcal{E}_j) \times \mathcal{E}_i$, choose $F^i(\cdot)$ so that the incentive constraints

$$\begin{aligned} \sum_{\omega \in E_i} p(\omega) u_i(f(\omega), \omega) &\geq \sum_{\omega \in E_i \text{ s.t. } \tilde{E}_i \cap E_{-i}(\omega) \neq \emptyset} p(\omega) u_i\left(f\left(\tilde{E}_i \cap E_{-i}(\omega)\right), \omega\right) \\ &+ \sum_{\omega \in E_i \text{ s.t. } \tilde{E}_i \cap E_{-i}(\omega) = \emptyset} p(\omega) u_i\left(F^i\left(\tilde{E}_i, (E_j(\omega))_{j \in N}\right), \omega\right) \end{aligned} \quad (3)$$

are satisfied. That is, $F^i\left(\tilde{E}_i, (E_j(\omega))_{j \in N}\right)$ gives the punishment for agent i when she says \tilde{E}_i at $E_i(\omega)$, and the remainder of agents say $(E_j(\omega))_{j \neq i}$. Any choice of $F^i(\cdot)$ such that these punishments are sufficiently great to make (3) hold is OK. To see that such a choice is always possible, fix the agent i , her information E_i and deviation \tilde{E}_i and note that one of the following three cases must hold:

1. $\tilde{E}_i \cap E_{-i}(\omega) = \emptyset$ for all $\omega \in E_i$: In this case, $F^i\left(\tilde{E}_i, (E_j(\omega))_{j \in N}\right) = a_f$ will satisfy constraint (3).
2. $\tilde{E}_i \cap E_{-i}(\hat{\omega}) \neq \emptyset$ for some $\hat{\omega} \in E_i$, and $i \in N_0(\hat{\omega})$ whenever $\omega \in E_i$ and $\tilde{E}_i \cap E_{-i}(\omega) = \emptyset$: So if $\omega \in E_i$ is such that $\tilde{E}_i \cap E_{-i}(\omega) \neq \emptyset$, and if $\omega' \in \tilde{E}_i \cap E_{-i}(\omega)$, then it follows that $\omega \rightarrow_i \omega'$. Thus whenever $\tilde{E}_i \cap E_{-i}(\omega) \neq \emptyset$, it holds that $f(\omega) = f\left(\tilde{E}_i \cap E_{-i}(\omega)\right)$. So setting $F\left(\tilde{E}_i, (E_j(\omega))_{j \in N}\right) = a_f$ whenever $\tilde{E}_i \cap E_{-i}(\omega) = \emptyset$ will satisfy the constraint.
3. $\tilde{E}_i \cap E_{-i}(\hat{\omega}) \neq \emptyset$ for some $\hat{\omega} \in E_i$, and there exists ω^* such that $\tilde{E}_i \cap E_{-i}(\omega^*) = \emptyset$ and $i \in N_\infty(\omega^*)$: Just choose $F^i\left(\tilde{E}_i, E_{-i}(\omega^*), E_i\right)$ such that $u_i\left(F^i\left(\tilde{E}_i, E_{-i}(\omega^*), E_i\right), \omega^*\right)$ is sufficiently low for (3) to hold. Such a choice is always possible since $i \in N_\infty(\omega^*)$.

Finally, construct $F : \times_{i \in N} \mathcal{E}_i \rightarrow A$ itself as follows. If $e = \left(\tilde{E}_i\right)_{i \in N} \in \times_{i \in N} \mathcal{E}_i$ is such that $\bigcap_{i \in N} \tilde{E}_i \neq \emptyset$, set $F(e) = f\left(\bigcap_{i \in N} \tilde{E}_i\right)$. (As above, this is a well defined construction). For all other points in $\times_{i \in N} \mathcal{E}_i$, define

$$A_e = \left\{ F^i(e, E_i) : i \in N, E_i \in \mathcal{E}_i \right\}$$

Then choose $F(e) = a_0$ where a_0 satisfies $u_i(a_0, \omega) \leq u_i(a, \omega)$ for all $i \in N, \omega \in \Omega$ and $a \in A_e$. Such a choice is always possible by Condition FD. By construction F satisfies (1) and (2), establishing that f is p -incentive compatible. Since $p \in P$ and f \mathcal{E}^* -measurable were chosen arbitrarily, this completes the proof.

Since \mathcal{E}^* is contractible information, it is clear that any coarser information (i.e. $\mathcal{G} \preceq \mathcal{E}^*$) is also contractible. I next prove the converse, namely that information is contractible *only if* it is a coarsening of \mathcal{E}^* . That is, no information finer than \mathcal{E}^* is contractible. Moreover, \mathcal{E}^* is the unique partition with this property. Thus \mathcal{E}^* represents the *maximum contractible information*.

Lemma 4 *A partition \mathcal{G} is contractible information only if $\mathcal{G} \preceq \mathcal{E}^*$.*

Proof. The proof is by contradiction. Suppose to the contrary that there exists some partition \mathcal{G} that is contractible information and for which $\mathcal{G} \not\preceq \mathcal{E}^*$. Now, \mathcal{E}^* not finer than \mathcal{G} is equivalent to⁶

$$\neg (\forall E \in \mathcal{E}^* \exists G \in \mathcal{G} \text{ s.t. } E \subset G)$$

or equivalently

$$\exists \hat{E} \in \mathcal{E}^* \text{ s.t. } \forall G \in \mathcal{G}, \neg (\hat{E} \subset G)$$

Take any such \hat{E} , along with $\omega_1, \omega_2 \in \hat{E}$ such that $\omega_1 \rightarrow \omega_2$ but $G(\omega_1) \neq G(\omega_2)$. (If no such pair of elements existed, we would have $G \in \mathcal{G}$ such that $\hat{E} \subset G$). So $\omega_1 \rightarrow_i \omega_2$ for some $i \in N$.

Define the function $f : \Omega \rightarrow A$ by

$$f(\omega) = \begin{cases} a_1 & \text{if } \omega \in G(\omega_1) \\ a_2 & \text{otherwise} \end{cases}$$

where $a_1, a_2 \in A$ are such that $u_i(a_1, \omega_1) < u_i(a_2, \omega_1)$. Such a choice is always possible by Condition SP. Since f is \mathcal{G} -measurable, by hypothesis it is p -incentive compatible for all $p \in P$.

By the revelation principle, there exists a direct mechanism F implementing f . So it must be the case that F is such that agent i prefers reporting $E_i(\omega_1)$ to $E_i(\omega_2)$ at $E_i(\omega_1)$:

$$\begin{aligned} \sum_{\omega \in E_i(\omega_1)} p(\omega) u_i(f(\omega), \omega) &\geq \sum_{\omega \in B} p(\omega) u_i(f(E_i(\omega_2) \cap E_{-i}(\omega)), \omega) \\ &+ \sum_{\omega \in E_i(\omega_1) \setminus B} p(\omega) u_i(F(E_i(\omega_2), (E_j(\omega))_{j \neq i}), \omega) \end{aligned} \tag{4}$$

where $B = \{\omega \in E_i(\omega_1) : E_i(\omega_2) \cap E_{-i}(\omega) \neq \emptyset\}$. Since $\omega_1 \rightarrow_i \omega_2$, $E_i(\omega_2) \cap E_{-i}(\omega_1) = E_N(\omega_2)$ so that (4) can be rewritten

$$\begin{aligned} p(\omega_1) (u_i(f(\omega_2), \omega_1) - u_i(f(\omega_1), \omega_1)) &\leq \\ \sum_{\omega \in B \setminus \{\omega_1\}} p(\omega) (u_i(f(\omega), \omega) - u_i(f(E_i(\omega_2) \cap E_{-i}(\omega)), \omega)) & \\ + \sum_{\omega \in E_i(\omega_1) \setminus B} p(\omega) (u_i(f(\omega), \omega) - u_i(F(E_i(\omega_2), E_{-i}(\omega)), \omega)) & \end{aligned} \tag{5}$$

Since $\omega_1 \rightarrow_i \omega_2$, it follows that $i \in N_0(\omega)$ for all $\omega \in E_i(\omega_1) \setminus B$. So define ε_1 and ε_2 by

$$\begin{aligned} \varepsilon_1 &= \max_{\omega \in B \setminus \{\omega_1\}} \{|u_i(a_1, \omega) - u_i(a_2, \omega)|\} \\ \varepsilon_2 &= \max_{\omega \in E_i(\omega_1) \setminus B} \{u_i(a_1, \omega) - z_i(\omega), u_i(a_1, \omega) - z_i(\omega)\} \end{aligned}$$

⁶ Here, \neg denotes the logical operator “not”.

The left hand side of (5) is strictly positive by construction, and is increasing in $p(\omega_1)$. On the other hand, the right hand side of (5) is bounded above by $(1 - p(\omega_1)) \max(\varepsilon_1, \varepsilon_2)$. So for all $p \in P$ with $p(\omega_1)$ sufficiently large, inequality (5) fails to hold, in contradiction to the hypothesis that f is p -incentive compatible. This completes the proof.

Thus we have:

Proposition 1 *The partition \mathcal{E}^* is the unique maximal contractible information partition.*

4 Examples

In the economy of Figure 1, we have $\omega_1 \rightarrow_1 \omega_2$ and also $\omega_1 \rightarrow_1 \omega_3$, and symmetrically $\omega_4 \rightarrow_1 \omega_3$. That is, agent 1 is always able to dishonestly claim to have the flu when he in fact has malaria, and to claim that he has a cold when in fact he is healthy. Also, $\omega_2 \rightarrow_1 \omega_3$ does not hold, since agent 2 would always recognize a dishonest claim by agent 1 to have a cold when in fact he has the flu. It follows that neither $\omega_2 \rightarrow_1 \omega_3$ or $\omega_3 \rightarrow_1 \omega_2$ holds.

To complete the analysis, we need to determine whether or not one of $\omega_2 \rightarrow_2 \omega_3$ or $\omega_3 \rightarrow_2 \omega_2$ holds. Certainly $\omega_2 \rightarrow_2 \omega_3$ and $\omega_3 \rightarrow_2 \omega_2$. Also, $\omega_2 \rightarrow_2 \omega_3$ only if $2 \in N_0(\omega_1)$ and $\omega_3 \rightarrow_2 \omega_2$ only if $2 \in N_0(\omega_4)$. That is, for agent 2 to be able to lie about whether she has observed any symptoms (i.e. whether she has observed $\{\omega_1, \omega_2\}$ or $\{\omega_3, \omega_4\}$) then it must be the case that she cannot be punished arbitrarily heavily when she lies and is found out. This happens if she denies agent 1 has symptoms when he in fact has malaria, of when she claims agent 1 has symptoms when in fact he is healthy.

Combining the above statements, we have $\mathcal{E}^* = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ if agent 2 can be punished arbitrarily heavily in both of the states ω_1 and ω_4 , and $\mathcal{E}^* = \{\Omega\}$ otherwise.

For the economy of Figure 2, it is straightforward to show that $\mathcal{E}^* = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ independent of the level of punishments available. This is an instance of the familiar result that since agent 1 and 2 can both fully distinguish $\{\omega_1, \omega_2\}$ from $\{\omega_3, \omega_4\}$, this information is contractible.

For a slightly more complex example, consider an economy with state space $\Omega = \{\omega_1, \omega_2, \dots, \omega_M\}$, with $M \geq N$. Each agent $i \in N$ is a specialist in identifying state ω_i , in the sense that $\mathcal{E}_i = \{\{\omega_i\}, \{\omega_j : j \neq i\}\}$. That is, agent i knows when state ω_i occurs, but otherwise cannot distinguish any of the other states. In such an economy, is the realization of the state available for contracting?

For any agent i , the information collectively possessed by all other agents is

$$\mathcal{E}_{-i} = \{\{\omega_j\} : j \in N \setminus \{i\}\} \cup \{\{\omega_j : j \in (M \setminus N) \cup \{i\}\}\}$$

Now, $\omega \rightarrow_i \omega'$ if and only if $\{\omega, \omega'\} \subset \{\omega_j : j \in (M \setminus N) \cup \{i\}\}$ or $\omega = \omega'$. Moreover, if $M > N$ then $\omega_i \rightarrow_i \omega'$ for any $\omega' \in \{\omega_j : j \in (M \setminus N)\}$, since $\omega_i \rightarrow_i \omega'$ and $E_i(\omega') \cap E_{-i}(\omega_i) = \{\omega_j : j \in (M \setminus N)\} \neq \emptyset$. So $\mathcal{E}^* = \{\Omega\}$ if $M > N$, and $\mathcal{E}^* = \{\{\omega_i\} : i \in N\}$ if $M = N$.

That is, if there is a specialist for *every* state in this economy ($M = N$), then collectively all *other* agents also know the information, and the specialist cannot lie. So in such a case all information is available for contracting.

However, introducing even one state that no-one is able to recognize (i.e. $M = N + 1$) radically alters this conclusion. In this case, for any state ω_i the coalition of agents $N \setminus \{i\}$ is never able to distinguish between ω_i and ω_M , which enables agent i to deny knowledge of state ω_i . It then follows that *no* information is available for contracting. This conclusion holds independently of the size of punishments possible.

5 Common knowledge and limited punishments

Contrary to what might have been expected, even in the two-person case maximum contractible information does not coincide with common knowledge.

Consider again the example of Figure 1. There, if punishments are unlimited in all states for both agents (i.e. if $N_\infty(\omega) = N$ for all $\omega \in \Omega$), then $\mathcal{E}^* = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ whereas $\mathcal{E}_1 \wedge \mathcal{E}_2 = \{\Omega\}$. That is, even though the agents have no non-trivial common knowledge, it is still possible to contract upon the employer’s information about the worker’s health, since the employer cannot misreport her information without a positive probability of detection.

On the other hand, if punishments are everywhere bounded, common knowledge and maximum contractible information do coincide in this example. This property is in fact general to two-agent economies, and can be generalized to the N -agent case as:

Proposition 2 *If $N_0(\omega) = N$ for all $\omega \in \Omega$, then $\mathcal{E}^* = \bigwedge_{i \in N} \mathcal{E}_{-i}$.*

Proof. The partition $\bigwedge_{i \in N} \mathcal{E}_{-i}$ is that induced by components of the graph $(\Omega, \leftrightarrow)$ where

$$\omega_1 \leftrightarrow \omega_2 \text{ if and only if } E_{-i}(\omega_1) = E_{-i}(\omega_2) \text{ for some } i \in N$$

If $\omega_1 \leftrightarrow \omega_2$ then $\omega_1 \rightarrow_i \omega_2$, and then $\omega_1 \rightarrow_i \omega_2$ since punishments are limited everywhere. Conversely, if $\omega_1 \rightarrow \omega_2$ then $\omega_1 \rightarrow_i \omega_2$ and thus $\omega_1 \rightarrow_i \omega_2$ for some agent i , so $\omega_2 \in E_{-i}(\omega_1)$ and hence $\omega_1 \leftrightarrow \omega_2$. Thus the graphs $(\Omega, \leftrightarrow)$ and (Ω, \rightarrow) have the same components, and so $\mathcal{E}^* = \bigwedge_{i \in N} \mathcal{E}_{-i}$, completing the proof.

The following corollary, which holds independently of the boundedness of punishments, is worth stating separately. The proof is simply the second half of the proof of Proposition 2.

Corollary 1 $\bigwedge_{i \in N} \mathcal{E}_{-i} \preceq \mathcal{E}^*$.

When applicable, this characterization greatly eases the application of Proposition 1. For instance, consider a three-agent economy in which agent 1 is a worker, agent 2 is a coworker and agent 3 is an employer. Suppose we are interested solely in agent 1’s effort, which has two components – hours worked, which can either

be long (L) or short (S), and intensity, which can either be (H) or low (0). So the state space is $\Omega = \{\omega_{LH}, \omega_{L0}, \omega_{SH}, \omega_{S0}\}$. Agent 1 knows exactly how hard he has worked, so $\mathcal{E}_1 = \{\{\omega_{LH}\}, \{\omega_{L0}\}, \{\omega_{SH}\}, \{\omega_{S0}\}\}$. The coworker, agent 2, observes only the intensity, so $\mathcal{E}_2 = \{\{\omega_{LH}, \omega_{SH}\}, \{\omega_{L0}, \omega_{S0}\}\}$, and the employer observes only hours worked, so $\mathcal{E}_3 = \{\{\omega_{LH}, \omega_{L0}\}, \{\omega_{SH}, \omega_{S0}\}\}$. What aspects of agent 1’s effort are contractible here? The answer follows easily from Corollary 1. Trivially $\mathcal{E}_{-2} = \mathcal{E}_{-3} = \mathcal{E}_1$. Moreover, $\mathcal{E}_{-1} = \mathcal{E}_1$. Thus $\bigwedge_{i \in N} \mathcal{E}_{-i} = \mathcal{E}_1$ – so all information is contractible. The fact that agents 2 and 3 separately see different and incomplete aspects of agent 1’s information does not prevent their joint information being useful.

6 Acquiring information

From the discussion so far, it should be clear that the availability or otherwise of unbounded punishments only affects the shape of the maximum contractible information in cases where deviations from truth-telling are sometimes but not always detected. Whenever the economy is such that deviations are either *always* or *never* detected, unbounded punishments play no role. In such cases, maximum contractible information should be again expected to coincide with the partition $\bigwedge_{i \in N} \mathcal{E}_{-i}$, and to reduce to common knowledge in the two-agent case. In this section I confirm this result, and then give an argument for why the case of always-or-never detected deviations can be expected to hold.

We start with a couple of definitions which formalize the property that deviations from truth-telling are either always or never detected. First,

Definition 3 *A pair of partitions $(\mathcal{G}_1, \mathcal{G}_2)$ is said to have the property of pairwise intersection if for any element of the common knowledge partition $\bar{G} \in \mathcal{G}_1 \wedge \mathcal{G}_2$ and any pair of partition elements $G_1 \in \mathcal{G}_1$ and $G_2 \in \mathcal{G}_2$ satisfying $G_1 \cup G_2 \subset \bar{G}$ it holds that $G_1 \cap G_2 \neq \emptyset$.*

Next,

Definition 4 *A set of partitions $\{\mathcal{G}_i\}_{i \in N}$ is said to satisfy global pairwise intersection (GPI) if \mathcal{G}_i and \mathcal{G}_{-i} satisfy pairwise intersection for all $i \in N$.*

When applied to information partitions, global pairwise intersection is exactly the property of single-agent deviations from truth-telling either being always or never detected when all others tell the truth. We then have

Proposition 3 *If $\{\mathcal{E}_i\}_{i \in N}$ satisfy GPI then the maximum contractible information \mathcal{E}^* is equal to $\bigwedge_{i \in N} \mathcal{E}_{-i}$.*

Proof. Observe that under GPI, $\omega_1 \rightarrow_i \omega_2$ holds if and only if $\omega_1 \rightarrow_i \omega_2$. “Only if” is immediate from the definitions. For the converse, suppose $\omega_1 \rightarrow_i \omega_2$. So $E_i(\omega_2) \cap E_{-i}(\omega_1) \neq \emptyset$, and so there exists $\bar{E} \in \mathcal{E}_i \wedge \mathcal{E}_{-i}$ such that $E_i(\omega_2) \cup E_{-i}(\omega_1) \subset \bar{E}$. So $E_i(\omega_1) \subset \bar{E}$, and thus $E_{-i}(\omega) \subset \bar{E}$ for all $\omega \in E_i(\omega_1)$. By GPI, $E_i(\omega_2) \cap E_{-i}(\omega) \neq \emptyset$ for all $\omega \in E_i(\omega_1)$, and so $\omega_1 \rightarrow_i \omega_2$.

The result then follows exactly as in the proof of Proposition 2.

Up to this point, the information \mathcal{E}_i of agents has simply been taken as given. Where does this information come from? Suppose that each agent i starts life with information given by \mathcal{E}_i^0 . The set of information partitions $\{\mathcal{E}_i^0\}$ satisfy GPI if agent j , based on her own information $E_j^0 \in \mathcal{E}_j^0$, is unable to rule out any combination of other agents information $\{E_i^0\}_{i \neq j}$.

GPI would seem a reasonable enough property for these “primal” information partitions $\{\mathcal{E}_i^0\}$, but over time agents may acquire new information. Without any loss of generality, this process can be thought of as occurring in a number of sequential rounds, $r = 1, \dots, R$. At each round agents may refine their information – regardless of the process by which this occurs, it is clear that the refinement must still be some coarsening of the total information possessed by all other agents. Thus if \mathcal{E}_i^r denotes the information of agent i at round r , that agent’s information at round $r + 1$ must be of the form

$$\mathcal{E}_i^{r+1} = \mathcal{E}_i^r \vee \mathcal{G}_i^r$$

where \mathcal{G}_i^r is some coarsening of the combined information of all other agents, \mathcal{E}_{-i}^r . Under this very general description of information acquisition, it can be shown that information after round $r + 1$ satisfies GPI whenever the information at round r did.

To establish this result, it is convenient to start by noting that:

Lemma 5 *Let $\mathcal{E}_1, \mathcal{E}_2, \mathcal{F}_1, \mathcal{F}_2$ be partitions of Ω . Then $\mathcal{E}_1 \wedge \mathcal{E}_2 \preceq (\mathcal{E}_1 \vee \mathcal{F}_1) \wedge (\mathcal{E}_2 \vee \mathcal{F}_2)$.*

Proof. $\mathcal{E}_l \preceq (\mathcal{E}_l \vee \mathcal{F}_l)$ for $l = 1, 2$. Thus

$$\mathcal{E}_1 \wedge \mathcal{E}_2 \preceq (\mathcal{E}_1 \vee \mathcal{F}_1) \wedge \mathcal{E}_2 \preceq (\mathcal{E}_1 \vee \mathcal{F}_1) \wedge (\mathcal{E}_2 \vee \mathcal{F}_2)$$

The key result is then:

Lemma 6 *Suppose that \mathcal{E}_1 and \mathcal{E}_2 are partitions of a finite set Ω and satisfy pairwise intersection. Then if \mathcal{F}_1 and \mathcal{F}_2 are coarsenings of \mathcal{E}_1 and \mathcal{E}_2 respectively, the pair of partitions $(\mathcal{E}_1 \vee \mathcal{F}_2, \mathcal{E}_2 \vee \mathcal{F}_1)$ also satisfies pairwise intersection.*

Proof. Suppose to the contrary that the pair of partitions $(\mathcal{E}_1 \vee \mathcal{F}_2, \mathcal{E}_2 \vee \mathcal{F}_1)$ does not satisfy pairwise intersection. So there exist $E_1 \in \mathcal{E}_1, E_2 \in \mathcal{E}_2, F_1 \in \mathcal{F}_1, F_2 \in \mathcal{F}_2$ and $\bar{F} \in (\mathcal{E}_1 \vee \mathcal{F}_2) \wedge (\mathcal{E}_2 \vee \mathcal{F}_1)$ such that $E_1 \cap F_2 \subset \bar{F}$ and $E_2 \cap F_1 \subset \bar{F}$ but $(E_1 \cap F_2) \cap (E_2 \cap F_1) = \emptyset$.

By Lemma 5 $\exists \bar{E} \in \mathcal{E}_1 \wedge \mathcal{E}_2$ such that $\bar{F} \subset \bar{E}$. Then certainly $E_1, E_2 \subset \bar{E}$, and so $E_1 \cap E_2 \neq \emptyset$ by the pairwise intersection property of \mathcal{E}_1 and \mathcal{E}_2 . Next, for $i = 1, 2$ write $F_i = \bigcup_{j=1}^{n_i} F_{ij}$ with $F_{ij} \in \mathcal{E}_i$ – this is always possible since \mathcal{F}_i is a coarsening of \mathcal{E}_i – and with $F_{1j} \cap E_2 \neq \emptyset$ and $F_{2j} \cap E_1 \neq \emptyset$. The elements F_{ij} are all contained in \bar{E} , and so by pairwise intersection again, $F_{1j} \cap F_2 = \bigcup_{k=1}^{n_2} (F_{1j} \cap F_{2k}) \neq \emptyset$, and thus $F_1 \cap F_2 = \bigcup_{j=1}^{n_1} (F_{1j} \cap F_2) \neq \emptyset$.

By hypothesis $(E_1 \cap E_2) \cap (F_1 \cap F_2) = \emptyset$. Define $\hat{F} = \bar{F} \cap (\bar{E} \setminus (F_1 \cap F_2))$. Now, $\hat{F} \neq \bar{F}$ since $F_1 \cap F_2 \neq \emptyset$ by above, and $\hat{F} \neq \emptyset$ since $E_1 \cap E_2$ is non-empty, has non-empty intersection with $\bar{F} \cap \bar{E}$, and by hypothesis is disjoint to $F_1 \cap F_2$.

To complete the proof, it is sufficient to establish that there exists a subset of \hat{F} that is a member of $(\mathcal{E}_1 \vee \mathcal{F}_2) \wedge (\mathcal{E}_2 \vee \mathcal{F}_1)$, which since \hat{F} is a *proper* subset of \bar{F} gives the required contradiction. Take any element G of $\mathcal{E}_1 \vee \mathcal{F}_2$. Then $G \cap (F_1 \cap F_2) \in \{\emptyset, G\}$ since $\mathcal{F}_1 \vee \mathcal{F}_2 \preceq \mathcal{E}_1 \vee \mathcal{F}_2$ and $G \cap (\bar{F} \cap \bar{E}) \in \{\emptyset, G\}$ since $(\mathcal{E}_1 \vee \mathcal{F}_2) \wedge (\mathcal{E}_2 \vee \mathcal{F}_1) \preceq \mathcal{E}_1 \vee \mathcal{F}_2$. Thus $G \cap \hat{F} \in \{\emptyset, G\}$. An identical argument establishes that for any $G \in \mathcal{E}_2 \vee \mathcal{F}_1$, $G \cap \hat{F} \in \{\emptyset, G\}$. Thus there must be some subset of \hat{F} that is a member of $(\mathcal{E}_1 \vee \mathcal{F}_2) \wedge (\mathcal{E}_2 \vee \mathcal{F}_1)$.

Applied to any agent i and coalition $N \setminus \{i\}$, Lemma 6 ensures that the pair of partitions $(\mathcal{E}_i^{r+1}, \mathcal{E}_{-i}^{r+1})$ satisfies pairwise intersection whenever $(\mathcal{E}_i^r, \mathcal{E}_{-i}^r)$ does. This gives:

Proposition 4 *The partitions $\{\mathcal{E}_i^{r+1}\}$ satisfy GPI provided $\{\mathcal{E}_i^r\}$ satisfies GPI. Thus information at round r satisfies GPI if the information at round 0 satisfies GPI.*

7 Unknown preferences

Up to this point, we have been concerned with characterizing what information an outside observer can conclude is always available for contracting, given that the observer knows all features of the economy *other than* the actual probabilities of different states (i.e. p). A closely related question is that of what information is known to be contractible by an observer who does not observe the *preferences* of agents,⁷ $\{u_i\}$ (though it is assumed that he does know the preferences satisfy Conditions SP and FD).

Now, if an outside observer knows *neither* the preferences $\{u_i\}$ *nor* the probability mapping p , then the maximal information that he can conclude is available for contracting is again $\bigwedge_{i \in N} \mathcal{E}_{-i}$. This observation is immediate from Proposition 2 and Corollary 1: If preferences are not known, then it is possible that punishments are limited for all agents in all states, and so the maximal contractible information is $\bigwedge_{i \in N} \mathcal{E}_{-i}$ – which we know is weakly coarser than the contractible information under any other assumption about the boundedness of punishments.

It is worth noting in passing that any $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable function f can be supported with a mechanism that is independent of the probability mapping p , and which depends on preferences only to the extent of making use of a punishment a_0 which is worse than any outcome in the range of f for all agents in all states (see the proof of Lemma 3).⁸ Thus not only can an outside observer conclude that $\bigwedge_{i \in N} \mathcal{E}_{-i}$ is contractible information even if he cannot observe preferences or probabilities, but a planner charged with designing a mechanism to support an

⁷ Arguably, preferences are easier for an outsider to observe than probabilities. Preferences are usually believed to be reasonably stable over time, allowing an observer to infer them from agents' past actions. On the other hand, if the world is non-stationary then historical data will be of little use in inferring the probabilities of future events.

⁸ This is of course just the familiar forcing contract, which punishes agents for inconsistent reports by imposing an outcome that all agents dislike.

$\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable function f can do so without any knowledge of the probabilities, and with only very limited knowledge of the preferences of agents.

What information could an outside observer who did observe probabilities *but not preferences* conclude is contractible? Certainly $\bigwedge_{i \in N} \mathcal{E}_{-i}$ will be available for contracting. Whether or not any additional information is available for contracting depends on the class from which preferences are being drawn.

On the one hand, if *no* restrictions are placed on the class of possible preferences (beyond conditions SP and FD) then one can show that $\bigwedge_{i \in N} \mathcal{E}_{-i}$ is the finest information partition that an outsider could conclude is definitely available for contracting. The proof is similar to that of Lemma 4.

On the other hand, a restriction on the preference class as mild as the requirement that $u_i(a, \cdot)$ be \mathcal{E}_i -measurable for all $a \in A$ is enough to lead to probability mappings p for which strictly more information than $\bigwedge_{i \in N} \mathcal{E}_{-i}$ is available for contracting even when preferences are arbitrary. To see this, consider the following simple example. The state space is $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, there are two agents $N = \{1, 2\}$, and their information is given by $\mathcal{E}_1 = \{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ and $\mathcal{E}_2 = \{\{\omega_2, \omega_3\}, \{\omega_4, \omega_1\}\}$. So

$$\bigwedge_{i \in N} \mathcal{E}_{-i} = \mathcal{E}_1 \wedge \mathcal{E}_2 = \{\Omega\}$$

However, for the probability mapping p defined by $p(\omega) = 1/4$ for all $\omega \in \Omega$, the information $\mathcal{F} \equiv \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_4\}\}$ is available for contracting. To see this, let $f : \Omega \rightarrow A$ be any \mathcal{F} -measurable function. We will show that f is p -incentive compatible.

Without loss, $f(\omega_1) = f(\omega_3) = a$ and $f(\omega_2) = f(\omega_4) = b$ for some $a, b \in A$. Take the truth-telling mechanism $\mathcal{M}_i = \mathcal{E}_i$ for $i = 1, 2$ and $F(E_1, E_2) = f(E_1 \cap E_2)$. Then agent 1's truth-telling constraint upon observing $\{\omega_1, \omega_2\}$ is

$$\frac{1}{2} (u_1(a, \{\omega_1, \omega_2\}) + u_1(b, \{\omega_1, \omega_2\})) \geq \frac{1}{2} (u_1(b, \{\omega_1, \omega_2\}) + u_1(a, \{\omega_1, \omega_2\}))$$

which is satisfied for any choice of a and b and any specification of preferences. A similar argument applies to all the other truth-telling constraints. Thus f is p -incentive compatible as claimed.

To summarize, when an outside observer does not know either the probability mapping p or the preferences of agents $\{u_i\}$, then the most information that he can conclude is definitely available for contracting is $\bigwedge_{i \in N} \mathcal{E}_{-i}$. However, if the outside observer knows *either* probabilities *or* preferences, then there are circumstances under which he can conclude that strictly more information than $\bigwedge_{i \in N} \mathcal{E}_{-i}$ is available for contracting.

8 Cross reporting

Thus far it has been established that maximum contractible information is a well-defined concept, and that for two important classes (bounded punishments, and information satisfying GPI) of economies is equal to the partition $\bigwedge_{i \in N} \mathcal{E}_{-i}$. This

finding may seem more intuitive under the following lattice theory result, the proof of which is in the appendix.

Lemma 7 $\bigwedge_{i \in N} \mathcal{E}_{-i} = \bigvee_{i \in N} (\mathcal{E}_i \wedge \mathcal{E}_{-i})$

Lemma 7 says that when punishments are bounded or information satisfies GPI, the maximum contractible information can be thought of as follows. Each agent i can misreport only that information that is not known by the coalition of all other agents. Thus agent i can only deviate from truth-telling when deviations respect the partition $\mathcal{E}_i \wedge \mathcal{E}_{-i}$. The maximum contractible information is then that obtained by combining these restrictions for all agents, resulting in the partition $\bigvee_{i \in N} (\mathcal{E}_i \wedge \mathcal{E}_{-i})$. By Lemma 7 and Propositions 2 and 3, this partition is precisely the maximum contractible information.

9 Relationship to literature

This is not the first paper to consider what information is easy to contract upon. In this section, I show how some previous results can be seen as special cases of the analysis conducted here.

9.1 Harris and Townsend (1981)

An early version of the result that any information shared by at least two agents is contractible can be found in Harris and Townsend (1981).⁹ In their framework, the state space is of the form

$$\Omega = \times_{m \in M} \Omega_m$$

with typical element $\omega = (\omega_1, \dots, \omega_M)$. Each agent i observes a subset M_i of the M random variables $\omega_1, \dots, \omega_M$. Thus the information of agent i is given by

$$\mathcal{E}_i = \{ \{ \omega : \omega_{M_i} = \theta_{M_i} \} : \theta_{M_i} \in \times_{m \in M_i} \Omega_m \}$$

It follows that if M_{-i} is defined by $M_{-i} = \cup_{j \neq i} M_j$, then the combined information of all agents other than i is given by

$$\mathcal{E}_{-i} = \{ \{ \omega : \omega_{M_{-i}} = \theta_{M_{-i}} \} : \theta_{M_{-i}} \in \times_{m \in M_{-i}} \Omega_m \}$$

The authors define the private information of agent i as those variables ω_m that are observed only by that agent. Denote these privately observed variables by M_i^P . Harris and Townsend's Theorem 2 establishes that any allocation satisfying the incentive constraints

$$\sum_{\tilde{\omega} : \tilde{\omega}_{M_i} = \omega_{M_i}} p(\tilde{\omega}) u_i(f(\tilde{\omega}), \tilde{\omega}) \geq \sum_{\tilde{\omega} : \tilde{\omega}_{M_i} = \omega_{M_i}} p(\tilde{\omega}) u_i(f(\delta_{M_i^P}, \omega_{-M_i^P}), \tilde{\omega})$$

⁹ See Theorem 2 and also Section 10.

$\forall i \in N, \omega \in \Omega, \delta_{M_i^P} \in \times_{m \in M_i^P} \Omega_m$ is incentive compatible i.e. the only incentive constraints that matter are those related to individuals' private information. An immediate implication is the result that any variable ω_m observed by at least two agents is freely contractible, since there are then *no* incentive constraints to satisfy. Formally, the set of "contractible" random variables is defined by $M^* = M \setminus \bigcup_{i \in N} M_i^P$, and any allocation that is measurable with respect to

$$\{\{\omega : \omega_{M^*} = \theta_{M^*}\} : \theta_{M^*} \in \times_{m \in M^*} \Omega_m\}$$

is incentive compatible. In light of Lemma 3, this implication is equivalent to simply noting that under the stochastic restriction imposed by Harris and Townsend, global pairwise intersection holds and thus

$$\mathcal{E}^* = \bigwedge_{i \in N} \mathcal{E}_{-i} = \{\{\omega : \omega_{M^*} = \theta_{M^*}\} : \theta_{M^*} \in \times_{m \in M^*} \Omega_m\}$$

9.2 Postlewaite and Schmeidler (1986)

In an implementation theory context, Postlewaite and Schmeidler (1986) define *non-exclusivity of information (NEI)* as holding if $E_{-i}(\omega) = E_N(\omega)$ for all states ω and all agents $i \in N$. They establish that under NEI incentive constraints place no restrictions on implementability. In terms of this paper, their result can be expressed as:

$$\bigwedge_{i \in N} \mathcal{E}_{-i} = \mathcal{E}_N \text{ if } E_{-i}(\omega) = E_N(\omega)$$

In fact,

Lemma 8 $\bigwedge_{i \in N} \mathcal{E}_{-i} = \mathcal{E}_N$ if and only if $E_{-i}(\omega) = E_N(\omega)$ for all $\omega \in \Omega$ and $i \in N$.

Proof. Suppose first that $E_{-i}(\omega) = E_N(\omega)$ for all ω and $i \in N$. So $\mathcal{E}_{-i} = \mathcal{E}_N$, and the result follows immediately.

Next suppose that $\bigwedge_{i \in N} \mathcal{E}_{-i} = \mathcal{E}_N$. Suppose that contrary to the hypothesis, there exists $\omega \in \Omega$ and $j \in N$ such that $E_{-j}(\omega) \neq E_N(\omega)$. Then

$$\bigwedge_{i \in N} \mathcal{E}_{-i} \preceq \mathcal{E}_{-j} \prec \mathcal{E}_N$$

since $E_{-j}(\omega) \neq E_N(\omega)$. But this contradicts $\mathcal{E}^* = \mathcal{E}_N$, and so completes the proof.

From Propositions 2 and 3 it then follows that if either punishments are limited for all agents (i.e. $N_0(\omega) = N$) or if $\{\mathcal{E}_i\}_{i \in N}$ satisfy GPI, then Postlewaite and Schmeidler's NEI condition is both sufficient *and* necessary for incentive constraints to place no restrictions on implementability.

9.3 The private core of a differential information economy

Yannelis (1991) defines the *private core* of a differential information economy¹⁰ as the class of resource-feasible allocations satisfying *private measurability* and the condition that no coalition of agents can redistribute its combined allocation in a privately measurable way while increasing the utility of all coalition members. Koutsougeras and Yannelis (1993) have shown that any allocation satisfying private measurability will satisfy coalitional incentive compatibility, and hence individual incentive compatibility. Thus any allocation in the private core is incentive compatible.

In the current paper, this property of the private core can be seen as follows. In the differential information economies considered in the above papers, the outcome set A is effectively a set of net-trade vectors, i.e. commodity transfers that sum to zero. For any mapping $f : \Omega \rightarrow A$, let f_j denote the net-trade made by agent j and f_{-j} the net-trade made by the coalition $N \setminus \{j\}$. The mapping f is said to be privately measurable if f_j is \mathcal{E}_j -measurable for every agent $j \in N$. Private measurability clearly implies that f_{-j} is \mathcal{E}_{-j} -measurable, and since $f_j + f_{-j} \equiv 0$ it also implies that f_j is \mathcal{E}_{-j} -measurable. Hence f_j is \mathcal{E}_{-i} measurable for all $i \in N$, and so is $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable. As this is true for any $j \in N$, the mapping f is also $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable, and hence incentive compatible by Lemma 3 and Corollary 1.

Although every allocation of the private core is $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable, the converse is clearly not true. For instance, the final example of Section 4 (with $N = M$) is an instance of a case in which $\bigwedge_{i \in N} \mathcal{E}_{-i} = \{\{\omega\} : \omega \in \Omega\}$ and so any mapping $f : \Omega \rightarrow A$ is $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable, but in which \mathcal{E}_j is strictly coarser than $\bigwedge_{i \in N} \mathcal{E}_{-i}$ for all agents j . Thus there exist $\bigwedge_{i \in N} \mathcal{E}_{-i}$ -measurable mappings f that cannot possibly lie in the private core.

Appendix: Omitted proofs

Proof of Lemma 7. First note that the order relations \preceq and \succeq induce a lattice over the set of all possible partitions of Ω . This lattice is isomorphic to that induced by set inclusion over the corresponding set of σ -algebras. It is a standard result that any lattice of sets is distributive (and hence modular as well).¹¹

The proof is by induction over N . The base case $N = 2$ is immediate.

¹⁰ See also Allen and Yannelis (2001) for a review of work on differential information economies.

¹¹ See, e.g., Davey and Priestley (1990). If a lattice is distributive then

$$\begin{aligned} a \wedge (b \vee c) &= (a \wedge b) \vee (a \wedge c) \\ a \vee (b \wedge c) &= (a \vee b) \wedge (a \vee c) \end{aligned}$$

for all lattice elements a, b, c . If a lattice is modular then whenever $a \geq c$

$$a \wedge (b \vee c) = (a \wedge b) \vee c$$

We are required to show that

$$\bigwedge_{i \in N} \bigvee_{j \neq i} \mathcal{E}_j = \bigvee_{i \in N} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i} \mathcal{E}_j \right)$$

Now, take any $i_0 \in N$. Then

$$\begin{aligned} \bigwedge_{i \in N} \bigvee_{j \neq i} \mathcal{E}_j &= \bigvee_{j \neq i_0} \mathcal{E}_j \wedge \left(\bigwedge_{i \neq i_0} \bigvee_{j \neq i} \mathcal{E}_j \right) \\ &= \bigvee_{j \neq i_0} \mathcal{E}_j \wedge \left(\mathcal{E}_{i_0} \vee \bigwedge_{i \neq i_0} \bigvee_{j \neq i, i_0} \mathcal{E}_j \right) \end{aligned}$$

by distributivity. Then by the inductive step

$$\bigwedge_{i \in N} \bigvee_{j \neq i} \mathcal{E}_j = \bigvee_{j \neq i_0} \mathcal{E}_j \wedge \left(\mathcal{E}_{i_0} \vee \bigvee_{i \neq i_0} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \right) \right)$$

Next note that

$$\bigvee_{j \neq i_0} \mathcal{E}_j \succcurlyeq \bigvee_{i \neq i_0} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \right)$$

so that by modularity

$$\begin{aligned} \bigwedge_{i \in N} \bigvee_{j \neq i} \mathcal{E}_j &= \left(\bigvee_{j \neq i_0} \mathcal{E}_j \wedge \mathcal{E}_{i_0} \right) \vee \bigvee_{i \neq i_0} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \right) \\ &= (\mathcal{E}_{i_0} \wedge \mathcal{E}_{-i_0}) \vee \bigvee_{i \neq i_0} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \right) \end{aligned} \quad (6)$$

Now

$$\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \preccurlyeq \mathcal{E}_i \wedge \bigvee_{j \neq i} \mathcal{E}_j = \mathcal{E}_i \wedge \mathcal{E}_{-i}$$

Thus since (6) is true for arbitrary $i_0 \in N$ then

$$\begin{aligned} \bigwedge_{i \in N} \bigvee_{j \neq i} \mathcal{E}_j &= \bigvee_{i_0 \in N} \left\{ (\mathcal{E}_{i_0} \wedge \mathcal{E}_{-i_0}) \vee \bigvee_{i \neq i_0} \left(\mathcal{E}_i \wedge \bigvee_{j \neq i, i_0} \mathcal{E}_j \right) \right\} \\ &= \bigvee_{i \in N} (\mathcal{E}_i \wedge \mathcal{E}_{-i}) \end{aligned}$$

■

References

- Allen, B., Yannelis, N.C.: Differential information economies: Introduction. *Economic Theory* **18**, 263–273 (2001)
- Arya, A., Glover, J., Young, R.: Virtual implementation in separable Bayesian environments using simple mechanisms. *Games and Economic Behavior* **9**, 127–138 (1995)
- Crémer, J., McLean, R.P.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* **53**, 345–362 (1985)
- Crémer, J., McLean, R.P.: Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* **56**, 1247–1257 (1988)
- Davey, B.A., Priestley, H.A.: *Introduction to lattices and order*. Cambridge, UK: Cambridge University Press 1990
- Duggan, J.: Virtual Bayesian implementation. *Econometrica* **65**, 1175–1199 (1997)
- Harris, M., Townsend, R.: Resource allocation under asymmetric information. *Econometrica* **49**, 33–64 (1981)
- Koutsougeras, L.C., Yannelis, N.C.: Incentive compatibility and information superiority of the core of an economy with differential information. *Economic Theory* **3**, 195–216 (1993)
- McAfee, R.P., Reny, P.J.: Correlated information and mechanism design. *Econometrica* **60**, 395–421 (1992)
- Maskin, E.: Nash equilibrium and welfare optimality. *Review of Economic Studies* **66**, 23–38 (1999)
- Matsushima, H.: Bayesian monotonicity with side payments. *Journal of Economic Theory* **59**, 107–121 (1993)
- Mirlees, J.A.: The theory of moral hazard and unobservable behaviour: Part I. *Review of Economic Studies* **66**, 3–21 (1999)
- Moore, J., Repullo, R.: Subgame perfect implementation. *Econometrica* **56**, 1191–1220 (1988)
- Osborne, M.J., Rubinstein, A.: *A course in game theory*. Cambridge, MA: The MIT Press 1994
- Postlewaite, A., Schmeidler, D.: Implementation in differential information economies. *Journal of Economic Theory* **39**, 14–33 (1986)
- Yannelis, N.C.: The core of an economy with differential information. *Economic Theory* **1**, 183–198 (1991)