Optimal Audit Policies with Correlated Types[*]

Patrick González[†]and Martin Boyer[‡]

November 24, 1999

Abstract

We propose a multi-agents adverse selection version of Townsend's [16] model of costly audits where the agents' types are correlated. Audits are used because agents have a limited ability to bear risk so that the Full Surplus Extraction (FSE) scheme la Crmer and McLean [6, 7] and McAfee and Reny [11] would be suboptimal here. It is shown that Townsend's result of an optimal marginal arbitrage between rent extraction and efficiency does not hold in the case of perfect correlation: FSE is feasible – even in dominant strategies – by devising a contract that put the agents in a prisoner's dilemma. A numerical simulation of the model is performed which suggests that the single agent model is not a good approximation of the multi-agents case.

JEL Classification: C63, D82.

Keywords: audits, asymmetric information, correlated information, costly state verification.

1   Introduction

Audits and various forms of costly monitoring are routinely performed in organizations to gather information that is usually available at no cost to

1

some of its members. These practices are usually rationalized by invoking that the organization performs better this way since revelation of information reduces adverse rent-seeking behavior by its members. Another argument is that revelation of information allows the Principal to secure for himself the economic surplus created in the organization.

Yet, it is well known today from the works of Crmer and McLean [6, 7] and McAfee and Reny [11][1] that, when the agents' information is correlated, efficient production and complete surplus extraction is achievable by the Principal using only sophisticated transfer schemes. If the Principal is only interested in maximizing the sum of the agents utilities or if all parties can contract before any agent actually learned their private information, another older result by d'Aspremont and Grard-Varet [9] shows that this is possible even if all information is disseminated independently across agents.

The question is then why is costly monitoring so pervasive? The cited results relies heavily on two assumptions:

- all parties are risk neutral;

- all parties have no limited liability.

Basically, the transfer schemes required to achieve efficiency with budget balance or total surplus extraction are highly variable and generally involve large positive and negative transfers. This would be socially costly with risk-averse parties and downright impossible (not credible) if parties have limited liability.

The moment when either one of these two assumptions is relaxed, information about private types becomes valuable so that one may optimally consider spending resources to gather it through costly audits. In this paper, we relax both assumptions; Robert [13] and Demougin & Gatvie [10] have shown that the efficient arrangement is then no longer implementable. We expect then the Principal to economize on cost by implementing costly audits.

The aim of the paper is to determine

- who should be audited;

- when specific audits should be performed;

---

[1]See Aoyagi [1] and Brusco [3] for recent references.

- how transfers should be conditioned on the agents voluntary announcements and on the results of performed audits.

We assume throughout that the organization is able to commit firmly to an audit policy procedure to avoid the renegotiation issue.

The rest of this paper is divided as follows. In the next section, we present a multi-agents version of Townsend's (1979) model of audits. In section 3, we consider first the polar cases of completely independent and of perfectly dependent types. In section 3.1, we develop an argument about symmetrical types spaces to get a tractable version of the model. The general case of imperfectly correlated types is explored with a numerical example in section 3.2. Section 4 concludes.

2 The Model

Consider an organization composed of $N + 1$ players. There is a Principal (player 0) facing a set $I = \{1, \ldots, N\}$ of N agents where each $i \in I$ is the "name" of some agent. Let $\mathcal{P}(I)$ be the power set of agents. Each agent i performs a task for the Principal from which he entails a private random cost (a type) of $\theta^i = \theta^i_k \in \Theta^i$ where the type value strictly increases with k. We let $\Theta = \times_{i \in I} \Theta^i$ be the set of $|\Theta| = J$ types profiles where $| \cdot |$ is the cardinal operator. We note $p(\theta)$ the probability that profile θ is realized. For simplicity, we assume that each agent has the same number T of possible types so that $J = T^N$. The type set is indexed with modulo T, that is if $T = 8$, then $\theta_{15} \equiv \theta_7$. Given a subset n of agents, the vector θ_n denotes their types profile; hence $\theta_I \equiv \theta$. To compensate him for this action, each agent shall be given a a wage w^i that may be contingent on the messages sent to the Principal about the agents' types and on the result of performed audits.

We assume that there is a constant return to scale technology in auditing. Auditing an agent then costs c to the Principal and reveals the agent's type. The Principal introduces audits by committing herself to an *audit policy.* An audit policy specifies how audits are to be carried on conditionally on the agents' reports. We will consider the case of *simultaneous* audits[2] and we define here an audit policy for that case.

---

[2]A readily extension to this model is to consider the case of *sequential* audits where the decision to audit an agent may depend of the result of auditing another agent.

By the Revelation Principle, a message from agent i to the Principal can be restricted to an element $m^i \in \Theta^i$. We let m denote a profile of messages. With simultaneous audits, an audit policy is a mapping from Θ into a distribution q over $\mathcal{P}(I)$; that is, with some probability q_n, only the subset n of agents, $0 \le |n| \le N$, will be audited. The set of audit policies (distributions) is the the unit simplex $\mathbb{S}^K$ in $\mathbb{R}_+^K$ where $K = |\mathcal{P}(I)| - 1 = 2^N - 1$. Given an audit policy and a profile of messages m, the Principal audits agent i with marginal probability

$$\sum_{n \in \mathcal{P}(I \setminus i)} q_{\{i\} \cup n}(m). \tag{1}$$

The objective of the Principal is to minimize the expected cost of compensating the agents and of pursuing a given audit policy. Given the wages w^i and the number $|n|$ of agents audited, that cost is

$$\sum_{i \in I} w^i + c|n|.$$

Agent i's ex post payoff is simply $u_i(w^i - \theta^i)$ where all u_i are strictly concave Von Neumann-Morgenstern utility functions with $\lim_{x \to 0} u_i'(x) = \infty$ to insure interior solutions.

Auditing the subset n of agents reveals surely their type profile θ_n. Given their messages m_n, we define an *audit result* to be a $|n|$-tuple a where each ordered element a_j of a is an integer from 0 to $T - 1$ that specifies by how many indexes the agent in n with the j^{th} name overstated his cost. A $|n|$-tuple of zeroes is then equivalent to say that all audited agents told the truth; when $|n|$ is large or unspecified, we note that result $a_{|n|}^0$. If $T = 8$ and $m_{\{2,4,5\}} = [\theta_3, \theta_5, \theta_2]$, then $a = (0, 2, 7)$ is an audit result that says that agent 2 told the truth, agent 4 lied by over-reporting his true cost by 2 and agent 5 lied by underreporting his true cost by 1 (that is, over-reporting by 7), since both agents 4 and 5 are of type θ_3. The set of possible audit results for any subset $n \ne \emptyset$ of agent is $A_{|n|}$ and has $|A_{|n|}| = T^n$ elements. If $n = \emptyset$, $|A_0|$ is defined to be 1 in the sense that the only new piece of information brought by performing the random audit policy was that no audits were actually performed.[3]

---

[3]That degree of formalization is made here for the sake of completeness. In section 3.1, we will invoke the classical argument that Nash implementation necessitates only to specify what happen in a truth-telling equilibrium and in single-agent defection from this equilibrium.

A wage to agent i contingent on the profile of messages m and on the information a produced by the audits is noted $w_{n,a}^i(m)$. When no audit are performed, we note the wage $w_\emptyset^i(m)$. Contingently on that event, that wage depends only the message m sent by the agents and their name i.

The organization is ruled by a contract that specifies the wages to be paid and the audit policy to be performed. We model the contracting process with a very standard sequential game:

1. The Principal hires the agents by offering them a contract; an agent that refuses the contract exit the game with a reservation payoff of $u_i(0)$ all normalized to zero and the Principal gets to make another take-it-or-leave-it offer to the remaining subset of agents (there is no discounting of the players' payoffs).

2. The agents learn their private costs θ^i that they will born by executing the task for the Principal.

3. All agents send simultaneously and privately a messages m_i about their private cost to the Principal.

4. According to the audit policy induced by m, random audits are performed upon a subset n of agents which reveal information a about n.

5. Each agent i accomplishes his task for which he receives a wage of $w_{n,a}^i(m)$. The agent's payoff is then $u_i(w_{n,a}^i(m) - \theta^i)$.

In the first stage, we assume that it is common knowledge that all agents always accept any contract that yields their reservation payoff. We assume that the social surplus associated to the accomplishment of the task is always positive ex post, whatever the agents' profile of types so that, even if an audit has been performed, the ex post production decision is never an issue.[4]

---

[4]A variant would be for each agent to execute his task in stage 3, that is prior the audits are performed. With non risk neutral agent, we would then have to specify agent i's payoff as $v(\theta^i, w^i)$. What matters, with respect to the equilibrium of the game, is how the agent's type affects his reporting behavior. With our formulation, wages and costs are perfect substitutes and an increase in the type always increases the marginal utility of a dollar in wage. In the variant, the wages and the costs will not be perfect substitutes and an increase in the type will have no effect on the marginal utility of a dollar if the function v is separable.

Refusing to execute the task in the last stage is never an option for any agent once he has accepted the contract. On the other hand, we assume that all agents have a limited liability so that their wage must be non negative (see Sappington, 1983). Hence, the lowest payoff an agent may get is $u_i(-\theta^i)$. This formulation would fit a situation where the task is fully contractible (so that no compliance may entail an arbitrarily large penalty) while the messages are imperfectly contractible: unless it is proven that he lied about his costs, an agent can always enforce the contractual wage. Wages are paid at the end of the game so that contracts randomizing the wage policy, contingent on the history of the game, will be dominated for risk averse agents.

Let us count the number of wages that must be specified following the announce of a message profile m by the agents. For each subset n of audited agents, there are $|A_{|n|}|$ possible results. The number of subsets in $\mathcal{P}(I)$ that have $k = |n|$ elements is $\binom{N}{k}$. It follows that, given any m, there are $1 + \sum_{k=1}^{N} \binom{N}{k} T^k = (T+1)^N = L$ possible contingencies we must take into account.[5] Let E be the set of these contingencies with typical element n, a. A contract is typically then composed on one hand of a function that maps Θ into $\mathbb{S}^K$ and, on the other hand, of N non negative wages functions $w^i(\cdot)$ that each maps $\Theta \times E$ into a non negative number $w_{n,a}^i(m)$ in $\mathbb{R}_+$ (the wage set).

We identify a contract with simultaneous audits by the vector of numbers it must specify which we note δ. We then let $D = (\mathbb{S}^K \times (\mathbb{R}_+)^{NL})^j$ be the set of contracts with simultaneous audits. A contract must specifies $J(K + NL)$ numbers. In the simplest many agents case (the *2×2-case*), that of two agents ($N = 2$) with two types ($T = 2$), to which we will often refer, we have $J = 4$, $K = 3$ and $L = 9$ so that each contract must specify $4(3 + 2 \cdot 9) = 84$ variables.

The final outcome of a contract depends on the agents' reporting behavior, the realization of the random variable $\tilde{\theta}$ and of the (unspecified) random variable governing the audit decision. In stage two though, the temporary outcome of a contract only depends on the realization of $\tilde{\theta}$, or more precisely, on the announce m made by the agents about θ. We note that temporary

---

[5]Another way to get that result is to augment the type set of each agent by a "null" type which represents the ex post type of an agent that has not been audited. Then, either the agent is audited, with T possible outcomes, or he is not audited and we say that he has the null type. There are then $(T+1)^N$ possible outcomes to the audits.

outcome $\delta(\theta)$ which is then a vector of $K + NL$ numbers that identify the distribution (K) of the random audit policy to be performed and the NL wages that may potentially be paid from that point on.

Assume that the agents report truthfully their types. The expected cost C of a contract δ to the Principal is then

$$\mathrm{E}(C(\delta)) = \sum_{\theta \in \Theta} \sum_{n \in \mathcal{P}(I)} p(\theta) q_n(\theta) \left(\sum_{i \in I} w_{n,a}^i(\theta) + c|n| \right) \tag{2}$$

In the 2×2-case, this becomes

$$\sum_{\theta \in \Theta} p(\theta)[q_{\{1\}}(\theta)(w_{\{1\},(0)}^1(\theta) + w_{\{1\},(0)}^2(\theta) + c)$$
$$+ q_{\{2\}}(\theta)(w_{\{2\},(0)}^1(\theta) + w_{\{2\},(0)}^2(\theta) + c) \tag{3}$$
$$+ q_I(\theta)(w_{I,(0,0)}^1(\theta) + w_{I,(0,0)}^2(\theta) + 2c) + q_\emptyset(\theta)(w_\emptyset^1(\theta) + w_\emptyset^2(\theta))],$$

where $q_\emptyset(\theta) = 1 - q_{\{1\}}(\theta) - q_{\{2\}}(\theta) - q_I(\theta) \geq 0$ so that q belongs to $\mathbb{S}^3$.

Under the same assumption, that is if $m = \theta$, the expected benefit of a contract δ to an agent i is

$$\mathrm{E}(U^i(\delta, \tilde{m})) = \mathrm{E}(U^i(\delta, \tilde{\theta})) = \sum_{\theta \in \Theta} \sum_{n \in \mathcal{P}(I)} p(\theta) q_n(\theta) u_i(w_{n,a_{|n|}^0}^i(\theta) - \theta^i). \tag{4}$$

2.1 Implementability

The fact that each agent's type is a private information to him, raises the issue of implementability of a contract. We will focus mainly on Bayesian-Nash implementation although our main result applies also in dominant strategy implementation. By the Revelation Principle, implementability is feasible if the contract satisfies the incentive compatibility constraints (IC) that induce an honest reporting behavior. With the usual notation, for any profile of type θ, let Θ_{-i} denotes the set of vector of types θ_{-i} where $\theta_{-i}^j = \theta^j$ for all $j \neq i$. Let

$$\mathrm{E}(U^i(\delta, \tilde{\theta})|\theta^i) = \sum_{\theta \in \Theta} \sum_{n \in \mathcal{P}(I)} p(\theta|\theta^i) q_n(\theta) u_i(w_{n,a}^i(\theta) - \theta^i),$$

denotes the expected utility of agent i, conditional on his type and given that he announces his true type $m^i = \theta^i$ and that he expect agent j to do the

7

same, that is, $\tilde{m}^j \equiv \tilde{\theta}^j$. With that notation, the expected utility of agent i if he lies while all other agents tell the truth so that, for any profile of type θ, some $\theta_{-i} \in \Theta_{-i}$ shall be reported is

$$E(U^i(\delta, \tilde{\theta}_{-i})|\theta^i) = \sum_{\theta \in \Theta} \sum_{n \in \mathcal{P}(I)} p(\theta|\theta^i) q_n(\theta_{-i}) u_i(w_{n,a}^i(\theta_{-i}) - \theta^i).$$

The *IC* constraints are then

$$E(U^i(\delta, \tilde{\theta}) - U^i(\delta, \tilde{\theta}_{-i})|\theta^i) \geq 0 \quad \forall \theta_{-i} \in \Theta_{-i}, \forall \theta \in \Theta, \forall i \in I. \qquad (IC)$$

The *IR* ex ante constraint is

$$E(U^i(\delta, \tilde{\theta})) \geq 0, \quad \forall \in I. \qquad (IR)$$

An implicit *IR* constraint is also assumed for the Principal; that is, the contract has a bounded value.

The set of optimal implementable contracts Δ_c that minimizes the expected cost of the Principal is then given by

$$\Delta_c = \operatorname{argmin}_\delta E(C(\delta)) \quad \text{s.t. } (IC) \text{ and } (IR).$$

We parameterize it by the cost c of performing a single audit. It is straightforward to show that Δ_c is a non-empty compact set.[6] When auditing is cost-less ($c = 0$) we are in effect in a case of complete information: we can assume that the Principal always commit to always audit both agents since that entails no loss of efficiency. Hence, $q(\theta) = 1$ and incentive compatibility is assured by giving a zero wage to all agents that lied. It is straightforward to see that first-best contracts in Δ_0 all set a wage $w_{n,a_{|n|}^0}^i(\theta) = \theta^i$ as to keep utility constant across states.

---

[6]The weak inequalities (IC) and (IR) describe a closed set of implementable contracts. We assume that $q_n(m)w_{n,a}^i(m) = 0$ for any event n of measure zero: since $q_n(m)w_{n,a}^i(m)$ is bounded, this implies that wages have an upper bound; hence, the set of implementable contracts is a compact set. Let δ^0, such that $w_{n,a}^i = m^i$ and $q_\emptyset = 1$ be the contract that maximize the agents surplus (they all announce θ^T). That contract satisfies (IC) and (IR) so that the set of implementable contracts is never empty.

3 Contracting with Costly Audits

When $c > 0$, auditing with certainty is generally inefficient because agents are risk averse: a sufficiently high probability that they will be audited can deter them from misreporting their type. Townsend (1979) made that point clear in the case of a single agent[7]. We first reproduce his result here in the case where there are many agents whose types are completely independent.

When types are independent, the information the Principal receives from one agent provides no information on any other agent. In this case, the optimal contract is such that there exist a cutoff cost $\overline{\theta}$ such that all agents that announced $\theta_k < \overline{\theta}$ are never audited while all that announced $\theta_k \geq \overline{\theta}$ are audited with some positive probability.

Because types are independent and the Principal is risk-neutral, there is no loss of generality in assuming that one agent's wage, or probability of being audited, depends only on his message and actual type (if he is audited); that is, the optimal contract will be separable in many independent contracts, one for each agent.

In such contract, an agent θ might be tempted to misreport his type in order to get, with probability $1 - Q$, a higher wage w than the wage w' he would get by telling the truth. If the agent is audited (with probability Q) and lied, though, he gets zero. His *IC* constraint with respect to such deviation is thus:

$$u(w' - \theta) \geq Qu(-\theta) + (1 - Q)u(w - \theta) \geq 0.$$

Setting that inequality to zero and solving for Q we can define the following function:

$$Q(w, w', \theta) = \frac{u(w - \theta) - u(w' - \theta)}{u(w - \theta) - u(-\theta)}. \tag{5}$$

The following lemma characterizes the derivatives of Q which we note Q_1, Q_2 and Q_3. The derivative of Q_1 with respect to w is noted Q_{11}.

Lemma 1. *i)* $Q_2 < 0$. *ii)* *When* $w \geq w' \geq 0$, $Q_1 \geq 0$ *and* $Q_{11} \leq 0$. *iii)* *When* $w \geq w' \geq 0$, $Q_3 \leq 0$.

---

[7]See Gale and Hellwig (1985), or Bond and Crocker [2] for a recent treatment.

Proof. The sign of the second derivative is obvious. To prove point ii), we derive the function

$$\phi(w) = \frac{f(w) - g}{f(w) - h}$$

twice where $f(w) = u(w - \theta)$, $g = u(w' - \theta)$ and $h = u(-\theta)$. Note that $f(w) \geq h$ and $g \geq h$. Then $\text{sign}(\phi'(w)) = \text{sign}((g - h)f'(w)) \geq 0$ and $\text{sign}(\phi''(w)) = \text{sign}((f - h)f''(w) - 2(f'(w))^2) \leq 0$.

iii) The case is trivial when $w = w'$. When $w > w'$, rewrite $Q(w, w', \theta)$ as a function $\Phi_{w'}$ of θ parameterized by w':

$$\Phi_{w'}(\theta) = \frac{f(\theta) - g_{w'}(\theta)}{f(\theta) - h(\theta)},$$

where $f(\theta) = u(w - \theta)$, $g_{w'}(\theta) = u(w' - \theta)$ and $h(\theta) = u(-\theta)$ and let $\phi_{w'}$ be the derivative of that function:

$$\phi_{w'}(\theta) = \Phi_{w'}(\theta) \cdot \left(\frac{f'(\theta) - g'_{w'}(\theta)}{f(\theta) - g_{w'}(\theta)} - \frac{f'(\theta) - h'(\theta)}{f(\theta) - h(\theta)} \right).$$

We want to show that $\phi_{w'}(\theta)$ is non-positive; this will be true if we can show that the term in the brackets is negative. For any given value of θ, the first ratio in the brackets decreases as w' is increased on $[0, w]$. Hence, we check the value of the difference in the brackets at its highest value as w' approaches zero: there the bracketed term takes the value zero. We conclude that $\phi_{w'}(\theta)$ is bounded above by zero. $\qquad\square$

We now characterize the optimal contract when types are independent.

Proposition 1. *The optimal contract specifies to audit an agent that announces θ_k with probability $q(\theta_k) = Q(w(\theta_k), w(\theta_1), \theta_1)$ where $w(\theta_k) = K + \theta_k$ is the optimal wage for $k = 2, \dots, T$. Hence q is a non decreasing function of θ with $q(\theta_1) = 0$.*

Proof. Fix the wages $w(\theta)$ so that they are non-decreasing in θ. Agents always consider misreporting their type in the hope of getting a higher wage. It follows that, for any agent of type θ_l, we must check the *IC* constraints only for types $\theta_k > \theta_l$. An agent of type θ_l will not misreport his type for θ_k as long as the probability of being audited when announcing type θ_k is no less than $Q(w(\theta_k), w(\theta_l), \theta_l)$. By lemma 1 and because $w(\theta_k) \geq w(\theta_l)$,

that number is less than $q(\theta_k)$ so the proposed contract does satisfy incentive compatibility. Clearly, once the $w(\theta)$ are fixed, the Principal will wish to set these probabilities as low as possible. The maximum value $Q(w(\theta_k), w(\theta_l), \theta_l)$ may take is at θ_1; it follows that the principal sets $q(\theta_k) = Q(w(\theta_k), w(\theta_1), \theta_1)$ for all θ_k.

For any fixed value of $w(\theta_1)$, the Principal then solves

$$v(w(\theta_1)) = \min_{w(\theta_2), w(\theta_3), \ldots} \sum_{\theta \in \Theta} p(\theta)(w(\theta) - Q(w(\theta), w(\theta_1), \theta_1)c)$$

$$\text{subject to} \quad \sum_{\theta \in \Theta} p(\theta)u(w(\theta) - \theta) \geq IR.$$

This is a convex program because Q was shown to be concave in its first argument in lemma 1. The necessary and sufficient first-order conditions yield

$$1 - u'(w(\theta) - \theta) \left(\frac{u(w(\theta_1) - \theta_1) - u(-\theta_1)}{(u(w(\theta) - \theta) - u(-\theta_1))^2} - \lambda \right) = 0 \quad \forall \theta \in \Theta \setminus \theta_1,$$

where λ is the Lagrange multiplier of the participation constraint. It follows that $w(\theta) - \theta$ must be equal to some constant $\kappa(w(\theta_1)) = K$ at the optimum.

In a second step, the principal minimizes $v(w(\theta_1))$ which is a convex function by the properties of minimum functions. By the theorem of the maximum, its derivative with respect to $w(1)$ is

$$p(\theta_1)(1 - \lambda u'(w(\theta_1) - \theta_1)) - \sum_{k=2}^{T} p(\theta_k)Q_2(w(\theta_k), w(\theta_1), \theta_1)c. \tag{6}$$

By lemma 1, the second term of (6) is always positive; we assume that the first term is such that (6) is positive.[8] In that case, $w(\theta_1)$ should be set at its lowest possible value 0. It follows that wages are non decreasing in θ as it was assumed in the beginning of the proof.

$\square$

---

[8]The multiplier must be positive since strengthening the IR constraint should increase the Principal's cost. If (6) take negative values, then the optimal $w(\theta_1)$ should be greater than zero. One must then check that the IC constraints of the higher types are not reversed toward that wage. We plan to refine that part of the proof in a future version of the paper.

In this setup where agent types are independent, the Principal is not able to use the message sent by one player to determine her strategy concerning the other player since such message has no informational value. Consider now the other extreme case where types are perfectly correlated. It is then easy to built a contract that yields an efficient outcome and that gives all the surplus to the Principal.

The idea is that if true, one agent's message about his type reveals the whole type profile. It follows that a profile of messages is either "consistent", in which case it implies a type profile with non-zero probability of occurrence; or "inconsistent", in which case we can be sure that at least one agents is lying. It suffices to commit to pay a zero wage whenever an inconsistent profile of messages is observed. Given our Bayesian-Nash assumptions about all agents' equilibrium expectations; that is, all agents choose their optimal report while expecting all other agents to tell the truth; the best thing any agent can do is to tell the truth.

Although audits are not even used in the previous set-up, they can be very effective when one is looking for dominant strategy implementation. Consider the case where agents may communicate among themselves prior to sending their message to the Principal; they may even verify each other type but we assume that they cannot sign binding colluding agreement. Clearly, the previous contract could miserably fail if we relax the assumption that all agents expect others to tell the truth: if an agent expect all other agents to lie, he might be better off to lie himself too. For instance, if types are perfectly positively correlated, a low-cost agent knows that the other agents have also a low cost and if he thinks that all other agents will lie by announcing a high cost, he would be better off to announce a high cost too to prevent getting zero. Dominant strategy implementation would require that each agent's payoff be independent of his message. For instance, with a pooling wage, each agent would tell the truth about his cost whatever his expectations about the other agents strategies. But a pooling wage is not efficient because the agents are risk averse to their cost.

We now show how audits can be used to get dominant strategy implementation with budget balance (the *IR* constraints are binding) and *efficiency*: the agents get their first-best payoff and *audits are never used in equilibrium*. Given a message profile m, let $C(m)$ be the partition of I such that all agents in $\chi \in C(m)$ sent consistent messages and all agents in $I \setminus \chi$ sent inconsistent messages to those sent by χ. We have a colluding equilibrium whenever the agents sent a profile of messages such that $I \in C(m)$. Note that $C(m)$ is

12

unique otherwise we would have at least two agents that belong to the same χ in the first partition and to a different χ' in another partition: but their messages must either be consistent or not. We then get the following result.

Proposition 2. *With perfectly dependent types and if an audit reveals whether an agent has lied or not, contracts that yield the first-best outcome (with no audit) $\delta^*(m)$ if $|C(m)| = 1$ and, otherwise, audit at least one individual per $\chi \in C(m)$ and pays*

$$ w^i_{\theta_n}(m) = \begin{cases} 0 & \text{if } m^i \in C(\theta_n); \\ \text{reward} & \text{else.} \end{cases} $$

implement the first-best allocation in dominant strategy if the reward > 0 *is sufficiently high.*

These contracts simply say that the Principal should pay the first-best wage if all messages are consistent, without any audit, and, otherwise, should perform a sufficient number of audits (enough to insure the truthfulness of all agents), at least one agent per subset $c \in C(m)$; pay all agents that lied zero and reward sufficiently those who told the truth as to make their *IC* constraints not binding. The idea is that, when type are perfectly dependent, any profile of messages in a colluding equilibrium that is altered by single message, that is, by one agent changing his announce, would result in a non-consistent profile of messages (note that telling the truth for all agents is one of these colluding equilibria). The contract then makes a strictly dominant strategy out of changing one's announce for the truth whatever the current colluding equilibrium: if some agents lie, I will be rewarded beyond whatever I would have got by participating in a colluding strategy; if all agents are telling the truth, I would be exposed and would get zero if I should lie because I would then create an inconsistent profile of messages.

The most striking feature of this contract is that its outcome does not depend on the audit cost c as long as this cost is finite because audits are never used in equilibrium.

3.1 Complexity

As it was illustrated in section 2, a contract quickly becomes a fairly complex object even for small values of the parameters T and N. This is of some concern because, we could always address the complexity issue with respect

13

to T by choosing a coarser measure of the type space: instead of having *very high, high, low* and *very low* cost, we could simply relabel types as *high* or *low*. No such relabeling seems natural for an increase in the number of agents. In this section, we run a series of argument to rationalize a much simplified model.

A way to simplify a contract is to simplify the probability space. We assume that all agents have the same preferences u, the same type set Θ^1 and by restricting the probability distribution over Θ to be symmetric in the sense that the probability of any state can be written as a function of the number of agents that are of each possible type in Θ^1. With such distribution, all agents are exactly all alike ex ante and differ only by their type ex post. We then restrict our search for an optimal contract to contracts that are symmetric with respects to the symmetries induced by our construction (that goal is rationalized in proposition 3 below).

By De Moivre's theorem[9], the number of T-tuples of non negative integers that sum to N is $S(N,T) = \binom{T+N-1}{N}$. The state space is then built as follows. Let $\mu : \Theta \to \mathbb{N}^T$ be the function that maps any type profile into the $1 \times T$ vector that specifies the number of agents that have type 1, 2, etc. Clearly, the image of μ has $S(N,T)$ elements. We specify a distribution f over the image $\mu_i(\Theta)$. Then we independently draw an element in the set of permutations of I which gives us a ranking r of the N agents and an element μ of $\mu_i(\Theta)$ according to the distribution f. ¿From μ we build an urn (a set) $B(\mu)$ where we put μ_1 balls labeled θ_1, μ_2 balls labeled θ_2, etc. By construction, there are N balls in $B(\mu)$. Finally, we draw successively the N balls out of $B(\mu)$ without replacement and we assign r_n the type labeled on the n^{th} ball drawn. That is, if $r_3 = 7$ and the third ball is labeled θ_4, then agent 7 has type θ_4 ($\theta^7 = \theta_4$).

The probability of observing a given type profile θ is then

$$p(\theta) = \frac{f\left(\mu_i(\theta)\right)}{\binom{N}{\mu_i(\theta)}},$$

where the denominator is the multinomial coefficient

$$\binom{N}{\mu_i(\theta)} = \frac{N!}{\Pi_{t=1}^T \mu_t!}.$$

[9]The references for the uncommented combinatorial arguments are [5] and [14].

That probability distribution function depends on θ through μ; it follows that if $\mu_i(\theta) = \mu_i(\theta')$, then $p(\theta) = p(\theta')$. This implies that all agents are completely symmetrical: like for the type for profile, the probability distribution of any vector $\theta_{n'}$ conditional on the realization of some vector θ_n depends only on $\mu_i(\theta_n)$. It is then straightforward to show that all agents have the same marginal type distribution. Under that construction, we say that all agents are *symmetric* ex ante.

Not only do we get a simplified state space with symmetry but a contract is likely to be simplified as well; that is, an optimal contract will pool many agents that are of an indistinguishable nature to the Principal. There are now only $S(N, T) < J$ distinct profiles of messages[10] that can be sent to the Principal since each message profile m information content is resumed by the reduced message profile $\mu_i(m)$. Beside, all agents that share the same type ex post are still identical with respect to their information set. We argue that there exists an optimal contract that treats all agents equally ex ante and all (announced) types equally at the interim stage, that is, before audits are performed.

Let I_{θ_k} be the subset of agents that declare being of type θ_k at the interim stage. We say that a contract is *symmetric* if the following conditions are satisfied: $\forall i, j \in I_{\theta_k}, \forall \theta_k \in \Theta^1, \forall n, n' \in \mathcal{P}(I)$ and $m, m' \in \Theta$ such that $\mu_i(m) = \mu_i(m')$ and $\mu_i(m_n) = \mu_i(m'_{n'})$; and $\forall a \in A(n)$,

$$q_n(m) = q_{n'}(m'),$$

and whenever $(i \in n) \wedge (j \in n') \wedge (\theta^i = \theta^j)$ or $(i \notin n) \wedge (j \notin n')$, then

$$w^i_{n,a}(m) = w^j_{n',a}(m').$$

Note that this implies that two agents i and j that have declared the same type have the same marginal probability of being audited. If $i, j \in I_{\theta_k}$, then $\mu_i(m_{\{i\} \cup n}) = \mu_i(m_{\{j\} \cup n}), \forall n \in \mathcal{P}(I \setminus i, j)$. Hence, $q_{\{i\} \cup n}(m) = q_{\{j\} \cup n}(m)$ for these n. Consider any n such that $j \in n$ but $i \notin n$; permuting j by i in n yields n' such that $\mu_i(m_{\{i\} \cup n}) = \mu_i(m_{\{j\} \cup n'})$ and $q_{\{i\} \cup n}(m) = q_{\{j\} \cup n'}(m)$.

---

[10]That account can be also be obtained as a special case of Polya's theorem: there are N identical agents that form a symmetric group S_N and each is to be painted of one of T colors. We distinguish colors (types) but not names. The total number of combinations is given by $|S_N|^{-1}(\sum_{g \in S_N} T^{\text{cyc}(g)})$ where $\text{cyc}(g)$ is the number of cycles in permutation g. That total amounts to $S(N, T)$.

These subsets come by pairs and exhaust the set of subsets to which i and j might respectively be joined to. Thus equation (1) yields the same sum.

The following proposition established that we can restrict our search of an optimal contract to the class of symmetrical contracts in that case.

Proposition 3. *If all agents are symmetric ex ante, then Δ_c includes a symmetric contract.*

Proof. Let $\delta \in \Delta_c$; we want to build a symmetrical contract $\delta' \in \Delta_c$ from δ. Consider agents i and j where agent i generates the less expected cost to the Principal (expected wage plus expected cost of monitoring). Let j's part of the contract be copied from i's part. That is, whenever $m^i = m^j$, let $q_{i \cup n}(m) = q_{j \cup n}(m), \forall n \in \mathcal{P}(I \setminus i, j)$ and let $w^i_{n,a}(m) = w^j_{n,a}(m), \forall n \in \mathcal{P}(I)$ whenever n and a imply either that i and j were not audited or that they overstated their costs in the same fashion. Implementability is unaffected by this change and the value of the contract cannot decrease since j now cost the same amount as i. Repeat that operation for the other agents to complete δ'. $\square$

Let μ be a reduced message profile; two agents that have sent the same message should have the same probability of being audited. An audit policy can be expressed as the probability that, for example, half of all agents that have declared type 1 and half of those that have declared type 3 will be audited.

Now, for two reduced profile μ^1 and μ^2 that are permutations to one another, we have to specify the same number of numbers for the audit policy. We can associate all these permuted profiles to a single partition η of N into T or less integers $\eta_1 + \eta_2 + \ldots + \eta_{|\eta|} = N$ which we represent as a $|\eta|$-tuple $\eta = (\eta_1, \eta_2, \ldots, \eta_{|\eta|})$ such that $1 \leq |\eta| \leq T$. Let $\mathcal{N}(N, T)$ be the set of these partitions; there are $|\mathcal{N}(N, T)|$ of them[11]. A partition η is an event that says that there were $|\eta|$ *kinds* of types announced: η_1 agents announced a type of the first kind, η_2 agents announced a type of the second kind; etc. If all agents send the same message then $\eta = (N)$, whatever that message was.

For each partition η, we must compute the number of reduced profiles of messages that are associated to it. This amounts to compute the number of distinct ways we can assign the elements of η to T types. Suppose first that

---

[11]That number is the coefficient of x^N in the series expansion of $\Pi_{k=1}^{T}(1-x)^{-k}$. Beside, if T is large, $|\mathcal{N}(N, T)|$ does not increase as T increases

16

$\eta_k \neq \eta_{k'}, \forall 1 \leq k, k' \leq |\eta|$. Then there T ways to assign a type to η_1, $T-1$ ways to assign a type to η_2, ... , $T+1-|\eta|$ ways to assign a type to $\eta_{|\eta|}$; for a total of $\Pi_{k=1}^{d(\eta}(T-k)$ ways to assign types to η. But that formula will lead to double counting if some elements of η are repeated. For instance, if $\eta = (1,1)$ and $T = 3$ then there are only three ways to having two agents spread into 2 equal formations of 1: one type must not be announced and there are three such candidates. We have T types that we must assign to the elements of η. Clearly, if there are repetition in the elements of η, e.g. $\eta_1 = \eta_2$ like in the preceding example, we should not count assignment of θ_1 to η_1 and θ_2 to η_2 as a distinct assignment than that of θ_1 to η_2 and θ_2 to η_1. For each repeated number of η, we need to divide by the number of indistinguishable permutations it generates. For instance, if 8 appears three times in η, than we must divide the permutations associated to 8 by 3!. Furthermore, $T - |\eta|$ are left out of θ; the same reasoning implies that the total should be pondered by $(T - |\eta|)!$. Suppose there are r_1 1's into η, r_2 2's, etc; and let $r(\eta)$ be the vector of the r_k's. There are thus $T - d(\delta)$ types that are left out which we count as a last element of r. Then, the total number of reduced message profiles associated to η is the multinomial coefficient $\binom{T}{r(\eta)}$.

Let π denotes the product of the elements of a t-tuple and let $\eta+1$ be the t-tuple such that one was added to each element of η. For each η we must specify $\pi(\eta) - 1$ numbers for an audit policy; that is, we can audit from 0 to $\eta_k \geq 0$ of the agents that have announced the k^{th} kind of type of η, times those of the k'^{th} kind of type, etc. For instance, if $N = 8$ and $\mu_1 = [0, 4, 1, 3]$ and $\mu_2 = [1, 3, 4, 0]$; then both are permutation of the partition $8 = 4 + 3 + 1$ that we note $\eta = (4, 3, 1)$ and $\pi(\eta+1) = 40$. So there are 40 configurations of types we must consider to audit and that requires an audit policy composed of $40 - 1 = 39$ numbers.

Hence, to compute the total number of numbers that must be specified by a symmetric audit policy, we first list all the partitions η of N and, for each of them, we count the number $\binom{T}{r(\eta)}(\pi(\eta+1) - 1)$. Summing these products gives us the true complexity of the symmetric audit policy:

$$\sum_{\eta \in \mathcal{N}(N,T)} \binom{T}{r(\eta)}(\pi(\eta+1) - 1). \tag{7}$$

For instance, if $N = 2$ and $T = 2$, the possible partitions are $\eta^1 = (2)$ and $\eta^2 = (1, 1)$. We then compute
$\binom{T}{r(\eta^1)}(\pi(\eta^1 + 1) - 1) + \binom{T}{r(\eta^2)}(\pi(\eta^2 + 1) - 1) = 2 \cdot (3 - 1) + 1 \cdot (4 - 1) = 7.$
We get a more complex case with $N = 5$ and $T = 3$. The possible partitions

are $\mathcal{N}(N, T) = \{(5), (4, 1), (3, 2), (3, 1, 1), (2, 2, 1)\}$. This yields

$$\sum_{\eta \in \mathcal{N}(N, T)} \binom{T}{r(\eta)} (\pi(\eta + 1) - 1) = 3(4 - 1) + 6(10 - 1)$$
$$+ 6(12 - 1) + 6(16 - 1) + 3(18 - 1) = 234,$$

numbers to be specified.

We now attempt to reduce the number of wages to be explicitly specified in a contract. Again, two permuted profiles μ^1 and μ^2 must specify the same number of wages so that we can work from η and sum over $\mathcal{N}(N, T)$. We will resume the contingency n, a with a single $(T + 1) \times |\eta|$ matrix α. The first row is the number of agents of each type that where not audited. The T subsequent rows $k = 2, \ldots, T + 1$ are the number of audited agents of each announced kind of type that have overstated their cost by k indexes. A matrix that has only zeros in the T lowest rows implies that no agent was audited. If there are some strictly positive integers only in the second row, then all audited agents told the truth. If the first row is zero, then all agents were audited. The sum of all elements of that matrix is N.

For example, if $N = 18$ and $T = 5$, we could have $\eta = (10, 5, 3)$ and

$$\alpha = \begin{pmatrix} 4 & 5 & 0 \\ 2 & 0 & 2 \\ 1 & 0 & 1 \\ 3 & 0 & 0 \end{pmatrix}.$$

This would read that five of the nine agents in the first kind of type of η were audited: 2 were telling the truth, one was overstating his cost by 1 and three by 2. The five agents in the second kind of type in η were not audited. All agents in the third kind of type of η were audited and 1 was lying by overstating his cost by 1. Given any partition η, there is a set $\mathcal{A}(\eta)$ of such matrices.

We can divide the agents into classes according to:

1. Thoses that declared being the k^{th} kind of type in η and were not audited; there are $\alpha_{1,k}$ of them per column k.

2. Thoses that declared being of the k^{th} kind of type in η, were audited and had their message confirmed; there are $\alpha_{2,k}$ of them per kind k.

3. Thoses that were audited and that lied; there are $\sum_{t=3}^{T+1} \alpha_{t,k}$ of them per column k.

Consider any optimal IC contract. Because the contract is IC, agents always tell the truth and these contingent wages are never paid since no agents ever find himself belonging to that class. It follows that these wages have no bearing neither on the Principal program nor on the (IR) constraints of any agents. Besides, if some of these wages are strictly positive, setting them all to zero only relaxes the (IC) constraints. Hence, without loss of generality, we can assume that agents in the third class always get zero.

It follows that we only need to specify a number of wages equal to the number of non zeros entries there are in the first two rows of α. Let $\kappa(\alpha)$ be that number. For all possible α given η We count $\kappa(\alpha)$ non-zero entries. There are $\binom{T}{r(\eta)}$ distinct reduced message profiles that yield η. Summing over $\mathcal{N}(N,T)$, we need to specify

$$\sum_{\eta \in \mathcal{N}(N,T)} \binom{T}{r(\eta)} \left(\sum_{j=1}^{|\mathcal{A}(\eta)|} \kappa(\alpha^j) \right),$$

wages.[12] The difficult part is to compute the second term in each product. We do the following. Given η, the total number of configurations of column k of α that can be made by partitioning η_k within $T+1$ rows is $S(\eta_k, T+1)$. Let $\sigma(\eta)$ be the vector of these numbers. The total number of distinct matrices α that can be made from η is $\pi(\sigma(\eta))$. A first approximation of $\kappa(\eta)$ would then be $2|\eta|\pi(\sigma(\eta))$ but many of these α have zeros in their first two rows for which no wage need to be specified. Now, if we knew how many zeros appears in all these matrices α, then we would know that a proportion $2(T+1)^{-1}$ of them appear in the first two rows and we could subtract these zeros. Let's first count how many times zero might appear if we rearrange column k in all possible fashions:

1. We may have from $\max(0, T+1-\eta_k)$ to T zeros in column k; pick z of them.

---

[12]Note that two distinct matrices $\alpha, \alpha' \in \mathcal{A}(\eta)$ that have the same first two rows must nevertheless specify different wages since they represent different events (the same number of agents of each kind of type lied but in a different fashion). Below, we argue that we can disregard these distinct events when we consider Bayesian-Nash implementation as all agents expect nothing but the other agents to tell the truth in these equilibria.

2. These z zeros may be disposed in $\binom{T+1}{z}$ ways in column k.

3. Once the z zeros have been disposed, there are $\binom{\eta_k - 1}{T - z}$ ways of disposing the η_k unlabeled units into the remaining $T + 1 - z$ labeled locations.

4. It follows that the number of zeros that will appear in column k is

$$\sum_{z=\max(0, T+1-\eta_k)}^{T} z \binom{T+1}{z} \binom{\eta_k - 1}{T - z}.$$

Now, each of these configurations of column k is to be matched with many different configurations of the other columns. If we do the exercise for all columns at once, then we find that there will be

$$Z(\eta) = \sum_{z_1 = \max(0, T+1-\eta_1)}^{T} (\ldots) \sum_{z_{|\eta|} = \max(0, T+1-\eta_{|\eta|})}^{T} \left(\sum_{k=1}^{|\eta|} z_k \right) \Pi_{k=1}^{|\eta|} \binom{T+1}{z_k} \binom{\eta_k - 1}{T - z_k},$$

zeros in all $\alpha \in \mathcal{A}(\eta)$.

The total number of wages that must be specified for a symmetric contracts is thus

$$\frac{2}{T+1} \sum_{\eta \in \mathcal{N}(N, T)} \binom{T}{r(\eta)} \left(|\eta|(T+1)\pi(\sigma(\eta)) - Z(\eta) \right). \tag{8}$$

In figure 1, we tabulated some of these totals for values of N and T up to 6. In each cell, the numerator is the sum of the total numbers of numbers needed for the optimal audit policy, equation (7), and that of those needed for the wages, equation (8). The denominator gives the same sum in the general case where the symmetric assumption is relaxed.

Finally, there is a another line of argument we can pursue to further tackle the curse of dimensionality. Because we focus on Nash implementation - all agents rationally expect the other agents to tell the truth in Bayesian-Nash equilibria - we do not need to specify wages for all contingencies that follows an audit. The only wages that matter are those that may happen in equilibrium or off-equilibrium given one-sided defection; that is, those where at least $N - 1$ agents are telling the truth.

Within a symmetric contract, this implies that for each reduced message profile μ (with associated partition η), and given that any audited agent

		T			
	2	**3**	**4**	**5**	**6**
N					
2	36.90% $\frac{7+24}{12+72}$	27.62% $\frac{15+72}{27+288}$	21.93% $\frac{26+160}{48+800}$	18.13% $\frac{40+300}{75+1\,800}$	15.43% $\frac{57+504}{108+3\,528}$
3	14.21% $\frac{16+84}{56+648}$	9.57% $\frac{46+468}{189+5\,184}$	7.28% $\frac{100+1\,680}{448+24\,000}$	5.91% $\frac{185+4\,650}{875+81\,000}$	4.98% $\frac{308+10\,836}{1\,512+222\,264}$
4	4.68% $\frac{30+224}{240+5\,184}$	2.73% $\frac{111+2\,184}{1\,215+82\,944}$	1.96% $\frac{295+12\,320}{3\,840+64\times10^4}$	1.55% $\frac{645+49\,600}{9\,375+324\times10^4}$	1.28% $\frac{1\,239+158\,928}{19\,440+12\,446\,784}$
5	1.39% $\frac{50+504}{992+38\,880}$	0.67% $\frac{231+8\,190}{7\,533+1\,244\,160}$	0.45% $\frac{736+70\,840}{31\,744+16\times10^6}$	0.34% $\frac{1\,876+409\,200}{96\,875+1\,215\times10^5}$	0.27% $\frac{4\,116+1\,787\,940}{241\,056+653\,456\,160}$
6	0.38% $\frac{77+1\,008}{4\,032+279\,936}$	0.15% $\frac{434+26\,208}{45\,927+17\,915\,904}$	0.09% $\frac{1\,632+340\,032}{258\,048+348\times10^6}$	0.06% $\frac{4\,795+2\,782\,560}{984\,375+4\,374\times10^6}$	0.05% $\frac{11\,914+16\,449\,048}{2\,939\,328+32\,934\,190\,464}$

Figure 1: Reduced complexity of an optimal symmetric contract.

		T						
		2	3	4	5	6	7	8
	2	23	51	90	140	201	273	356
	3	56	172	388	735	1 244	1 946	2 872
	4	110	447	1 255	2 845	5 607	10 010	16 602
N	5	190	987	3 376	9 026	20 496	41 426	76 728
	6	301	1 946	7 968	24 815	64 330	146 160	300 612
	7	448	3 528	17 040	61 160	179 544	455 652	1 035 336
	8	636	5 994	33 726	138 215	456 183	1 285 767	3 212 583

Figure 2: Reduced complexity of an optimal Bayesian-Nash symmetric contract.

that lied gets zero, we only need to specify a different sent of wages for each configuration of audits: audit results do not matter since all agents expect nothing but the truth in the equilibria we are interested in. More precisely, we only need α matrices composed of two rows: those that were audited those that were not. There are $\eta_k + 1$ ways of auditing agents that have declared being of the k^{th} kind of type. If this was the only reported type, we would only need to specify $2(\eta_k - 1)$ wages but if there is a second kind of type k' that can be audited in $\eta_{k'} + 1$ ways, we need $2(\eta_k - 1)(\eta_{k'} + 1)$ wages for kind k and $2(\eta_{k'} - 1)(\eta_k + 1)$ wages for kind k'. Hence, the number of wages to be specified for η is

$$2\pi(\eta + 1) \sum_{k=1}^{|\eta|} \frac{\eta_k - 1}{\eta_k + 1}.$$

This reduces the total number of numbers to be specified for a symmetric contract in Bayesian-Nash implementation to

$$\sum_{\eta \in \mathcal{N}(N,T)} \binom{T}{r(\eta)} \left(\left(\pi(\eta + 1) \left(1 + 2 \sum_{k=1}^{|\eta|} \frac{\eta_k - 1}{\eta_k + 1} \right) \right) - 1 \right). \tag{9}$$

Some of these numbers are tabulated in figure 2. Hence, there are only 23 variables to specify in the 2×2-case. In the next section, we proceed to a numerical resolution of that case in the space of probability matrices. Those 23 variables are:

1. the probability of auditing both agents when both report a low cost;

2. the probability of auditing a single agent when both report a low cost;

3. the probability of auditing both agents when both report a high cost;

4. the probability of auditing a single agent when both report a high cost;

5. the probability of auditing both agents when agents report different costs;

6. the probability of auditing only the low cost agent when agents report different costs;

7. the probability of auditing only the high cost agent when agents report different costs;

8. the wage to be paid when both agents report a low cost and are audited;

9. the wage to be paid when both agents report a low cost and are not audited;

10. the wage to be paid to the audited agent when both agents report a low cost and only one of them is audited;

11. the wage to be paid to the agent who is not audited when both agents report a low cost and only one of them is audited;

12. the wage to be paid when both agents report a high cost and are audited;

13. the wage to be paid when both agents report a high cost and are not audited;

14. the wage to be paid to the audited agent when both agents report a high cost and only one of them is audited;

15. the wage to be paid to the agent who is not audited when both agents report a high cost and only one of them is audited;

16. the wage to be paid to the low cost agent when agents reports different costs and are both audited;

23

17. the wage to be paid to the high cost agent when agents reports different costs and are both audited;

18. the wage to be paid to the (audited) low cost agent when agents reports different costs and only the low cost agent is audited;

19. the wage to be paid to the (not audited) high cost agent

when agents reports different costs and only the low cost agent is audited;

20. the wage to be paid to the (not audited) low cost agent when agents reports different costs and only the high cost agent is audited;

21. the wage to be paid to the (audited) high cost agent when agents reports different

costs and only the high cost agent is audited;

22. the wage to be paid to the low cost agent when agents reports different costs and are not audited;

23. the wage to be paid to the high cost agent when agents reports different costs and are not audited.

3.2   Contracting with imperfectly correlated types

The dominant strategy contract described in proposition 2 will also work with imperfectly correlated types: all that is needed is to commit to audit whenever some "suspect" profile of types is announced and to reward the whistle blower if some fraud is revealed afterward with a bonus at least as high as what he would have gain by participating to the conspiracy. Such a contract will be IC but then all profiles of types will have a positive probability of occurrence so that there will be a positive probability of auditing. At the margin, the Principal might want to economize on the expected cost of auditing by reducing the probability of auditing.

To make sense of the case of imperfectly correlated types, we need a good parameterization of probability space. We propose such parameterization for the 2×2-case in figure 3. For any given Θ^1, we need to specify the probability distribution f over

$$\mu_i(\Theta^1) = \{(2 \text{ types } \theta_1), (1 \text{ type } \theta_1 \text{ and } 1 \text{ type } \theta_2), (2 \text{ types } \theta_2)\}.$$

We represent that distribution in the projection of the simplex $\mathbb{S}^3$ in the $X \times Y$ space where each point (x, y) corresponds to a probability matrix

$$P(x, y) = \begin{bmatrix} x & \dfrac{1 - x - y}{2} \\ \dfrac{1 - x - y}{2} & y \end{bmatrix}.$$

In figure 3, the three vertex matches the subset of matrices for which types are perfectly dependent is the hypotenuse plus the origin (in thick lines). We have already established that one can achieve the FSE allocation in these cases. The set of matrices for which Townsend's result shall apply is given by the dotted downward slopping curve going from $(0, 1)$ to $(1, 0)$. We will

24

$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$

$p>1/2$
$\begin{bmatrix} p & 0 \\ 0 & 1-p \end{bmatrix}$

$\begin{bmatrix} q/2 & (1-q)/2 \\ (1-q)/2 & q/2 \end{bmatrix}$

$\begin{bmatrix} 1/2 & 0 \\ 0 & 1/2 \end{bmatrix}$

$\begin{bmatrix} y & \frac{1-y}{2} \\ \frac{1-y}{2} & 0 \end{bmatrix}$

$q>1/2$

$p<1/2$
$\begin{bmatrix} p & 0 \\ 0 & 1-p \end{bmatrix}$

$1/4$

$\begin{bmatrix} \frac{1}{4} & \frac{1}{4} \\ \frac{1}{4} & \frac{1}{4} \end{bmatrix}$

$q<1/2$

$\begin{bmatrix} \tau^2 & \tau_i(1-\tau) \\ \tau_i(1-\tau) & (1-\tau)^2 \end{bmatrix}$

$\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$

$1/4$

$\begin{bmatrix} 0 & 1/2 \\ 1/2 & 0 \end{bmatrix}$

$\begin{bmatrix} 0 & (1-x)/2 \\ (1-x)/2 & x \end{bmatrix}$

Figure 3: Space of symmetric probability matrices in the 2×2-case.

25

refer to that curve as Townsend's ridge for geometrical reasons that will appear clear below. These are the matrices for which agent i's knowledge of his own type does not inform him in any way on agent j's type; that is,

$$y + \frac{1 - x - y}{2} = \frac{y}{y + \frac{1 - x - y}{2}}.$$

Solving this implicit relation for $x + y \leq 1$ yields the equation of that curve

$$y = 1 + x - 2\sqrt{x}.$$

These matrices can also be represented by the parametric form $\begin{bmatrix} \tau^2 & \tau_i(1-\tau) \\ \tau_i(1-\tau) & (1-\tau)^2 \end{bmatrix}$ where τ runs from 0 to 1.

At this stage, we present a numerical experiment we did with the following parameterization:

$$u(z) = \log(11 + z),$$
$$IR = \log(11)$$
$$\Theta^1 = \{1, 10\},$$

We have computed the optimal contract δ on a triangular grid

$$x \text{ from } 0.025 \text{ to } 0.975,$$
$$y \text{ from } 0.025 \text{ to } 1\text{-x}.$$

over the projection of $\mathbb{S}^3$. Not surprisingly, when both agents report a low cost, no audit is ever performed. In fact, no audit is ever performed on an agent that declares low cost. Yet, as the probability of facing a high cost agent is increased, the rise in the low cost's wage one can expect in Townsend's model seems to be a local feature, or a "ridge" (see figure 4). An almost identical pattern emerges for the wage to be paid to the low cost agent when messages are mixed and the high cost agent is audited (we observe small differences along the x-axis next to the origin).

When both agents report a high cost, they are usually both audited when (x, y) is on the left side of the ridge although there is a mixed result when y is very low (see figure 5). No single audit is ever performed in that case

26

Figure 4: Wage when both agents report low cost (Townsend's ridge).

Figure 5: Probability of a dual audit when both agents report high cost. That probability goes from 1 to zero along the Townsend's ridge.

neither. When reported types are mixed, we obtain the opposite solution: the high cost agent seems to be only audited when we are on the right side of the ridge (see figure 6). There seems to be a transition along the ridge as the probability of no audit (instead of a single audit of the high cost agent) increases when the probability of getting a high cost increases.

4   Conclusion

We developed a multi-agents representation of Townsend's (1979) model of optimal audits. All our results tend to show that the results one can gather with the single agent representation are absolutely not typical.

In our understanding, the limit result of proposition 2 of a contract that implements the first-best allocation in dominant strategy with budget balance when types are perfectly dependent is a clear indication that the assumption that states are not verifiable is not a good approximation of a world where states are difficult and costly to verify.

References

[1] Masaki Aoyagi. Correlated types and bayesian incentive compatible mechanisms with budget balance. *Journal of Economic Theory*, 79:142–151, 1998.

[2] E.W. Bond and K.J. Crocker. Hardball and the soft touch: The economics of optimal insurance contracts with costly state verification and endogenous monitoring. *Journal of Public Economics*, 63:239–254, 97.

[3] Sandro Brusco. Unique implementation of the full surplus extraction outcome with correlated types. *Journal of Economic Theory*, 80:185–200, 1998.

[4] J. Caballé and A. Pomansky. Mixed risk aversion. *Journal of Economic Theory*, 71:485–513, 1996.

[5] Daniel I. A. Cohen. *Basic Techniques of Combinatorial Theory*. John Wiley & Sons, New York, 1978.

/h1/econo1/prof/pgon/articles/aca/qM_S2

Figure 6: Two views of the probability of a single audit of the high cost agent when agents report mixed costs. That probability goes from zero to one along the Townsend's ridge.

[6] J. Crémer and R.P. McLean. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica*, 53(2):345–361, 1985.

[7] J. Crémer and R.P. McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica*, 56:1247–1257, 1988.

[8] K. Dachraoui, G. Dionne, L. Eeckhoudt, and P. Godfroid. Proper risk behavior. unpublished.

[9] C. d'Aspremont and L. A. Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[10] Dominique M. Demougin and Devon A. Garvie. Contractual design with correlatad information under limited liability. *Rand Journal of Economics*, 22(4):477–489, Winter 1991.

[11] R.P. McAffe and P.J. Reny. Correlated information and mechanism design. *Econometrica*, 60:395–442, 1992.

[12] John Moore. Implementation, contracts, and renegotiation in environments with complete information. In Jean-Jacques Laffont, editor, *Advances in Economic Theory, Sixth World Congress*, number 20 in Monograph, pages 182–282, Great-Britain, 1992. Econometric Society, Cambridge University Press.

[13] Jacques Robert. Continuity in auction design. *Journal of Economic Theory*, 55:169–179, 1991.

[14] Fred S. Roberts. *Applied Combinatorics*. Prentice-Hall, Englewood Cliffs, New Jersey, 1984.

[15] David E. M. Sappington. Limited liability contracts betwen principal and agent. *Journal of Economic Theory*, 29:1–21, 1983.

[16] R.M. Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory*, 21:265–293, 1979.