Repeated Contracting in Decentralized Markets*

Sambuddha Ghosh
Boston University

Seungjin Han
McMaster University

April 1, 2013

Abstract

This paper studies a model where multiple principals repeatedly offer short-term contracts to three or more agents, each privately informed about her type. Agents observe contracts and actions, but principals observe agents' private messages only. We propose a simple class of mechanisms that is sufficient to sustain all equilibrium allocations in the repeated game when discounting is low. An equivalence theorem shows how only direct mechansims may be used to compute a principal's minmax value relative to arbitrarily general mechanisms. Endogenous monitoring by agents allows weaker notions of incentive compatibility than one-shot contracting, lowering players' mimnax values and supporting more equilibrium payoffs.

KEYWORDS: relational contracts, repeated contracting, competing mechanisms, folk theorems, endogenous monitoring

*We thank Madhav Chandrasekher, Ani Dasgupta, Joyee Deb, Larry Epstein, Yuk-fai Fong, Ying Lei, Maxim Ivanov, Bart Lipman, Dilip Mookherjee, Andy Newman, Juan Ortner, Frank Page, Mike Peters, Edward Schlee, Caixia Shen, Andy Skrzypacz, and Lixin Ye for comments. Seminar audiences at Arizona State University, Boston University, Indiana University, Sydney University, University of Guelph, University of New South Wales gave useful suggestions. This research was initiated in 2010 during Han's sabbatical at Boston University; Han thanks the Department of Economics at Boston University for the hospitality. Errors remain our responsibility.

Contents

1 Introduction

Many economic problems observed in practice have both repeated contracting, followed by trading governed by the prevailing contract. Often, contracting is decentralized— several principals offer competing contracts rather than one grand contract that governs all the terms of trades in the market. Contracting is often non-exclusive. For example, buyers nonexclusively purchases goods from multiple sellers, either for final consumption or as an intermediate good. Sellers seldom, if ever, contract with buyers on terms of trades for the entire future, either because full committment is not possible or future contingencies are too complicated to specify exhaustively. Instead, sellers (principals) repeatedly offer trading mechanisms to buyers (agents) to negotiate the terms of trade for short periods of time.

The relational contracts literature discussed later shows how repeated interaction, usually in conjunction with short-term contracts, replaces long-term contracts that specify terms of trade for the entire future. However, this literature limits itself to models with only one principal. Many applications are more naturally modelled as *several* principals interacting repeatedly with several agents, the latter with private information that is payoff-relevant for the agents themselves and possibly even for the principals. At each interaction, all principals offer decentralized contracts that govern actions for the current period. Agents observe their type, the mechanisms offered, and send private messages to the principals who then take actions[1] according to their mechanisms and received messages. The latter observe nothing other than private messages from agents. This paper provides a tractable framework for such problems, without recourse to *ad hoc* restrictions on permissible mechanisms.

Before we study the dynamic model proposed above, let us see why competition among principals leads to complexity even in the static version. Why can sellers not ask buyers to report just their types, as in direct mechanisms ubiquitous in models with a single principal? Since buyers have information on not only their types but also all sellers' mechanisms, a seller can usually improve his payoff by asking buyers to also report the mechanisms offered by the other sellers, so that he may best respond to them. This logic repeats *ad infinitum,* with each seller asking buyers to disclose how the other sellers will react to information about his mechanism, and so on; this leads to what is sometimes called the "infinite regress problem". Indeed, Epstein and Peters (1999) shows that the resulting class of universal mechanisms resembles the set of hierarchies of beliefs; we refer to these as "complex mechanisms".

We prove that any strictly individually rational payoff profile (i.e. each player gets above his/her minmax) derived from incentive compatible one-shot mechanisms may be supported

[1] In our example with multiple sellers and buyers, the action can be interpreted as a profile of quantity and payment pairs, one for each buyer. Generally, an action is any decision that a principal and agents can contract on; it includes the signing of employment contracts with a subset of applicants, specifying compensation based on the profits of the employer and other performance evaluation measures; an action can also specify a task or an effort for agents, together with an incentive scheme to induce it.

in a perfect Bayesian equilibrium (henceforth PBE) of the repeated contracting game. By itself this is not particularly useful for applications: in the static setting, one cannot compute minmax values because these are defined by taking minimum and maximum over complex mechanisms, for which there is no parsimonious expression; in addition, equilibrium short-term mechanisms offered in the literature are much more involved than direct mechanisms (DMs) because a deviation by a principal must be detected and punished immediately.

Our results lend themselves to applications for two reasons. First, our equivalence theorem provides a simple algorithm to calculate minmax values with respect to the set of *all* complex mechanisms; it shows that the complex minmax of any principal in the repeated game can be expressed as his *maxmin* value when he can offer only actions, and the other principals are restricted to *DMs*. Since actions and DMs are both much simpler objects than complex mechanisms, the calculation of the complex minmax reduces to a programming problem. No algorithm exists for calculating the corresponding minmax value in static games, even in finite ones.

Second, our sufficient mechanisms are also simple. While punishments must be included in the contract in the static game, it can be put into continuation payoffs in the dynamic game when players are patient. We show that on the equilibrium path, principals only need to offer deviator-reporting direct mechanisms (DDMs) at each time. DDMs are only slightly more complicated than DMs in that agents report their types, as well as the player (if any) who must be punished from the next period. DDMs assign current actions contingent on agents's type reports only; reports on who deviated influence the continuation play only. Therefore, DDMs can be viewed as DMs augmented by agents' cheap talk messages about the identity of the deviating player, if any. Off the path following a deviation by principal j, our equivalence results show that punishments can be meted out with extended direct mechanisms (EDMs) — in addition to the information requested by a DDM, principals $k \neq j$ ask agents to report which action principal j would play if all agents were to follow the prescribed strategies.[2]

Our results are particularly simple in the special case of private values, where each agent's type influences only her own utility. We show that DDMs are enough both on and off the path, obviating the need for EDMs. Furthermore, various definitions of minmax value are equivalent. Numerical examples illustrate the computational ease of our equivalence theorems.

It is interesting to see how endogenous monitoring by agents in the dynamic setting expands the set of equilibrium payoffs. First, endogenous monitoring allows players to completely neutralize a deviating principal's ability to make his action choice contingent on agents' types off the path following his deviation. In other words, if principal j is being punished, agents induce the same action from the the principal's mechanism regardless of their types. Consequently, j cannot benefit by offering a complex mechanism. If agents observe some

[2]While EDMs are required off the path, the principal's minmax value relative to complex mechanisms is expressed in terms of direct mechanisms.

other action realized from the principal's mechanism, they can infer that some agent deviated from the messages she was supposed to send. Because agents can report such a deviation to principals in order to punish the deviating agent in the following periods, such a deviation can be deterred. This is why a principal's minmax can be calculated as if he can only offer actions instead of complex mechanisms.

Secondly, endogenous monitoring weakens the notion of incentive compatibility (IC). As we will show later in details, some false type reports result in unexpected actions with positve probability, and will be subsequently reported by agents. In the dynamic setting, contemporaneous incentive compatibility is not required to deter such false reports. Deferring the punishment to the continuation game permits more mechanisms on the equilibrium path and, at the same time, allows more stringent punishments, forcing minmax values in the dynamic setting below those in the static setting.

We review the related literature in the next subsection. Section 2 provides an example that illustrates the key idea behind our results. Section 3 formally sets up the model for dynamic competing mechanisms. Section 4 presents the results for the general case of interdependent values. Section 5 shows how results simplify under private values. Section 6 concludes the paper by discussing contributions and possible generalisations.

1.1 Relation to Literature

An extensive lierature on relational contracts (see among others Pearce and Stachetti (1999), Fong and Li (2010), and Levin (2003)) shows how repeated interaction provides incentives for short-term contracts instead of contracts that specifies terms of trade for the entire future; however this literature restricts us to a single principal, whereas most real world problems involve multiple principals offering competing mechanisms. In addition the main concern of this literature is show that repeated contracting provides a way to provide incentives even when some key variables are non-contractible. Our results is most relevant when everything is contractible, and the model is liable to be intractably complex; fortunately our results salvage the situation by identifying a simple and sufficient class of mechanisms.

McAfee (1993) and Peck (1995) pointed out that the standard revelation principle defined over agents' payoff types fails when multiple principals (e.g., sellers) compete in designing trading mechanisms. Menu theorems by Peters (2001) and Martimort and Stole (2002) show that when there is only one agent in the model, an equilibrium in complex mechanisms may be replaced by each principal offering a menu of contracts.[3] In the static model wtih three or more agents, Yamashita (2010) shows that each principal can offer mechanisms where each agent reports not only her type but also recommends the DMs that principals should offer,

[3] Han (2006) showed that menu theorems extend to multiple agency problems if contracting is bilateral, i.e principals negotiate separate contracts with each agent. Prat and Rustichini (2003) is an example of bilateral contracting.

covering all possible off-equilibrium situations there is no unanimous recommendation. In all of the above, minmax values are couched in terms of the complex mechanisms; since there is no closed-form expression for the latter, it is not easy to identify which payoffs can be supported in equilibrium. This and other connections are further explored after the model is introduced.

The incentive-compatible extended DMs in Pavan and Calzolari (2009) generate only a subset of the equilibrium allocations possible under complex mechanisms. Attar, Mariotti and Salanié (2011) studies a market with adverse selection as a non-exclusive common agency problem.[4] Our work is closer to Bergemann and Välimäki (2003), which characterises truthful Markov perfect equilibrium payoffs of a dynamic common agency with symmetric information. They restrict principals are restricted to offering nonlinear reward schedules to the agent, effectively ruling out richer mechanisms *a priori*: We *derive* the optimality of simple mechanisms.

The folk theorem in Peters and Trancoso Valverde (2013) characterizes equilibrium allocations in a one-shot setting without distinguishing between principals and agents, because everyone can offer mechanisms and send messages (over two rounds). Our setting is standard: agents and principals are distinct, and agents send a single round of messages to principals; this is important as direct communication among principals is often forbidden as collusion. Our algorithm derives minmax values even under incomplete information, whereas individiual rationality in the above paper translates into minmax values under complete information but not under incomplete information.

In the literature on contract theory, it is of important to design a contract that ensures proper monitoring among agents (e.g., Ma (1989), Obara and Rahman (2010), Chandrasekher (2012)). Our paper shows that monitoring is particularly simple in the dynamic setting because agents simply report the indentity of a deviating player; such a report is cheap talk as far as the current period is concerned because it does not affect the current action choice: Monitoring leads to evolution of the continuation equilibrium.

One of our contributions is to draw attention to the complexity of contracts. The literature on incomplete contracting takes it as given that contracts cannot always specify all contingencies, and examines its implications. Segal (1999) shows that when the environment becomes arbitrarily complicated, the optimal mechanism is a very simple one. Our results suggest that a contracting environment with multiple principals, while potentially very complex, could very well involve simple equilibrium mechanisms. Evans (2008) investigates a general model of hold-up with renegotiation and finds simple efficient contracts.

[4] They demonstrate how to extend their results in the common agency framework to the biltareal contracting framework.

2 An Application

We present here an applicaton to illustrate our result in a simple setting, and postpone until the end a discussion of how various restrictive assumptions used in this example may be relaxed.

2.1 Imperfect Substitutes and Bertrand Competition

There are two manufacturers, 1 and 2, each producing at constant marginal cost c. Products are close but imperfect substitutes. Manufacturers repeatedly sell their products through I retailers with $I \geq 3$. Each manufacturer j sets the price of its product and also faces a binary choice of marketing effort $e_j \in \{0, 1\}$ such as advertisement, either directly, or through local retailers, with the manufacturer absorbing the marketing cost r. Retailers are assumed to compete to sell each firm's product, leaving them with zero (supernormal) profits.

Retail markets operate a la Bertrand, but with differentiated products.[5] Specifically, if manufacturer i's price is p_i and manufacturer j's price is p_j, then i sells $Q^s_{i\varepsilon} = A_i s_i - 2p_i + p_j + \varepsilon$ units. The last term ε is the underlying uncertainty in the market and it follows a distribution H with mean zero. Neither manufacturers nor retailers observe the realization of ε. We let $Q^s_i := \mathbb{E}_\varepsilon(Q^s_{i\varepsilon}) = A_i s_i - 2p_i + p_j$, where $\mathbb{E}_\varepsilon(\cdot)$ is the expectation operator over ϵ.

Each parameter $s_i \in [\underline{s}, \overline{s}]$ follows a distribution F independently and A_i depends on two factors — 1) market conditions, as respresented by a state in $\{G, B\}$, and 2) the marketing effort profile $e = (e_1, e_2)$ of the two manufacturers. The element of the state-space $\Theta = [\underline{s}, \overline{s}]^2 \times \{G, B\}$ is common knowledge among retailers but not known to manufacturers.

In state (s, G) the market is good enough that an advertisement by either manufacturer results in the demand for both products going up; in state (s, B) an advertisement by i destroys some of the demand of j, possibly securing some of it for product j. In the interest of simplicity, we make the extreme assumption that, in state B, i can destroy j's maket completely, but gains none of the lost customers; hence

$$A_i = \begin{cases} 1 + \max\{e_1, e_2\} & ; \theta = G \\ 1 - e_j & ; \theta = B \end{cases}.$$

Both G and B are equally likely. If manufacturer i chooses $e_i = 1$, he pays a cost r that is so large that it reduces net profits. This ensures that it is never profitable to advertise to increase demand. Assume that manufacturers and retailers have reservation profits equal to zero. Both assumptions simplify calculations and are not critical to the results.

In each period, a manufacture can write one-period contracts with its retailers, and commit to take certain actions as a function of messages from the retailers. In principle, there are

[5] Alternatively, we can consider retail markets that operate *a la* Cournot.

no restrictions on the complexity of the message space and what information it can contain. The manufacturer does not observe the competing manufacturer's mechanism, her terms of trade with retailers, let alone the state in the retail market; indeed he observes only the sales of his product and does not directly observe the price set by the other manufacturer. Due to underlying uncertainty on the demand shock in the market and a lack of "deep pockets", retailers may not commit to pay a fixed amount of payment to the manufacturer. Therefore, the outcome of a manufacturer's contract for a retailer is a pair comprising the price of the product that the manufacturer charges to the retailer and the retailer's payment, which is the manufacturer's price multiplied by the quantity that the retailer actually sells in the retail market. Assume that sales of each product are distributed equally over all retailers who charge the lowest price. If retailers purchase the product from the manufacturer at the same price, competition among retailers restricts them to zero profits[6]; manufacturers maximise the discounted sum of profits using a common discount factor δ.

We first find the joint-profit maximising solution. In state $s \in \{s_1, s_2\}$, let $(\hat{p}_1^s, \hat{p}_2^s, \hat{e}_1^s, \hat{e}_2^s) \equiv (\hat{p}^s; \hat{e}^s)$ denote the joint-optimal price and marketing levels:

$$(\hat{p}_1^s, \hat{p}_2^s, \hat{e}_1^s, \hat{e}_2^s) \in \underset{(p^s;e^s)}{\arg\max} \sum_{i=1,2} \left\{ Q_i^s(p_i^s - c_i) - re_i^s \right\}.$$

Since r is large, the optimal marketing effort level is zero in every state: $\hat{e}_1^s = \hat{e}_2^s = 0$. The first-order conditions are: $\hat{p}^s$ solves

$$\begin{array}{c} \frac{\partial Q_1^s}{\partial p_1^s}(p_1^s - c_1) + Q_1^s + \frac{\partial Q_2^s}{\partial p_1^s}(p_2^s - c_1) = 0 \\ \frac{\partial Q_2^s}{\partial p_2^s}(p_2^s - c_2) + Q_2^s + \frac{\partial Q_1^s}{\partial p_2^s}(p_1^s - c_2) = 0. \end{array}$$

This reduces to

$$4p_1^s - 2p_2^s = s_1 + c_1; \; -2p_1^s + 4p_2^s = s_2 + c_2.$$

Henceforth assume $s_1 = s_2 = s$, and $c_1 = c_2 = c$. Then we have

$$\hat{p}_1^s = \hat{p}_2^s = (s + c)/2.$$

Second order conditions for a maximum hold. Therefore, the maximum expected joint profit is $\hat{\pi} = .5\mathbb{E}_s(s - c)^2 > 0$, where $\mathbb{E}_s(\cdot)$ is the expectation operator over s.

[6]This is common, for example when there is an MRP, a maximum retail price. We see this as a simplifying assumption. We can accomodate more complicated models where the retailer makers a fixed industry-standard mark-up, or even engages in Nash bargaining with the manufacturer. This is not central to our story and hence kept as uncomplicated as possible.

Direct Mechanisms and the One-shot Game

Can manufacturers sustain production levels and marketing effort levels that maximize the joint surplus without complex contracting in this repeated trading and contracting environment? The answer is in the affirmative. Each manufacturer j then offers a direct mechanism (DM) π_j^d in which each retailer i is asked to report the true state. If more than half of retailers report the same state s, then the contract (i.e., quantity supplied and payment) assigns to each retailer i a price-payment pair

$$\left(\hat{p}_j^s, q_{ij}^s \hat{p}_j^s\right) \text{ for all } s \in \{s_1, s_2\},$$

where $\hat{p}_j^s$ is the joint profit-maximizing price and it is the price that manfacturer j charges to each retailer and q_{ij}^s is the quantity that retailer i buys from manufacturer j, which eventually the quantity that the retailers will sell in the retail market. If manufacturers maintain their DMs, they share the joint surplus generated in the retail market as above depending on the value of s; before learning the value of s the expected profit of each manufacturer is $\hat{\pi}/2$.

However we now show that if the game has only one period, the joint-optimal quantities cannot be sustained with the DMs (π_1^d, π_2^d) above, even when these are augmented by cheap talk on the identity of a deviator. Note that cheap talk plays no role in a one-shot setting. The best response function p_i^* of manufacturer i is:

$$p_i^*(p_j) = (s + p_j)/4 + c/2,$$

and therefore $(s + c)/2$ is not a best response to $(s + c)/2$; since

$$\left.\frac{\partial \pi_i}{\partial p_i}\right|_{p_i = p_j = (s+c)/2;\ e_i = e_j = 0} < 0,$$

each manufacturer wants to lower prices. The intution is clear — lowering p_i raises i's demand but lowers the demand for j, but only the first effect is considered in calculating the best response whereas both effects interact in determining the joint optimum. Manufacturer 2 would rather price his product lower; manufacturer 1 can respond to this by offering a contract where he asks each reatiler if manufacturer 2 deviated and choosing a different price-payment pair if relaiers report a deviation. We refer to this as manufacturer 1's first order contract, to distinguish it from a DM, which is a zero-order contract because it is based only on the market information. Manufacturer 2 could then ask retailers if 1 offered a retaliatory contract and in turn can best respond to this, giving us a second order contract. This sequence of best responses leads to the infinte regress problem of Epstein and Peters (1999) that was discussed earlier: the sufficient class of contracts becomes an intractable object and there is no way to compute the minmax payoffs of the principals.

Sustaining Collusion in the Repeated Contracting Game

In the repeated game, we can however sustain the above DMs with cheap talk. To do so manufacturer j augments the direct mechanism π_j^d by also asking retailers to report any player who unilaterally deviated: This report has no effect on the current allocation and hence is purely cheap talk for the current period; however it influences future play. A DM augmented by such a report is called a DDM, for deviator-reporting direct mechanism. The DDM bases future play on the report that is delivered by a majority of agents. As long as each manufacturer j offers the direct mechanism π_j^d, every retailer reports the true state because $I \geq 3$ ensures that a unilateral false report does not affect the majority report; on the equilibrium path she reports 0 as the deviator, to mean that there are no unilateral deviations. We simplifed the model in such a way that agents cannot gain by reporting falsely. We need to check that principals do not have profitable deviations from the above DDMs. If j deviated from the DDM π_j^d augmented by the cheap talk, he is reported to the other manufacturer by all retailers and subsequently minmaxed for a certain number of periods.[7] Proposition 3 will show that the above DDMs can be sustained in equilibrium if we can show that the equilibrium payoff $\hat{\pi}/2$ of each principal exceeds his minmax value, provided the principals are sufficiently patient. The lower the minmax value the lower the threshold discount factor for which co-operation can be sustained.

We leave a detailed examination of the dynamic incentives for later, and calculate minmax values of a principal, say 2; since our example is symmetric both principals have the same value. First consider the case where 1 tries to punish 2 without using any information from the agents (the retailers); we refer to this as punishing with simple actions. Irrespective of s, the worst punishment 1 can inflict on 2 is by choosing $p_1 = 0 = e_1$ so that $Q_2 = s_2 - 2p_2$ is lowest, whatever p_2 is. Given this, firm 2's best response is to offer a direct mechanism that asks agents to report the value of s and chooses

$$p_2^s \in \arg\max_{p_2^s} \, (s - 2p_2^s)(p_2^s - c),$$

which gives a price of $p_2^s = (s + 2c)/4$ and a profit of

$$\underline{\pi} = \int [\max_{p_2^s}(s - 2p_2^s)(p_2^s - c)]dF(s) = \frac{1}{8}\mathbb{E}_s(s^2 - 4c^2), \tag{1}$$

assumed to be strictly positive. Therefore, the minmax value without using any information from the agents is $\underline{\pi}$.

[7] If manufaturer j deviates from the DDM π_j^d, the price of its product would change. This in turn would change the retailers' choices of quantities from the non-deviating manufacturer. However, manufacturer k cannot infer whether manufacturer j deviated or not just by observing the quantities demanded by retailers. This is because market demand for k's product also depends on the underlying uncertainty ε in the market, which is not observable by anyone. This is why retailers need to report deviations.

However manufacturer 1 can mete out a harsher punishment by committing to a direct mechanism in which he asks all retailers to report the state and chooses $e_1^B = 1$ if the majority of retailers report (s, B) and $p_1^G = 0 = e_1^G$ if the majority reports (s, G). Manufacturer 2 is instructed to offer a DDM, where he asks retailers to report the state and if manufacturer 1 deviated. This leads to a profit of $\underline{\pi}/2$, which is the minmax value based on information from the agents.

Hence asking the agents to report information on the best way of punishing the deviating principal makes it easier to punish and hence to sustain collusion.

2.2 Comments on the Example

We note the main messages of the example, and discuss the special features that simplified the analysis.

1. The first observation is that simple actions are not enough to mete out severe punishments. In this example the most severe punishment requires using a DDMs, i.e. direct mechanisms augmented by a report on the identity of unilateral deviators.

2. This raises a moot question, Is this true in general that DDMs are enough? Our theorems show that DDMs are enough on the equilibrium path. In this example they suffice even off the path beacuse incentive problems are absent when it comes to extracting information for purposes of devising punishemnts for a deviating principal. To relax this, we must provide incentives to agents to reveal the truth, which in turn requires agents to be patient. In our example that all information is common knowledge among agents; this means that the above-mentioned incentive problem is avoided and the logic works even if the agents are myopic; as a result DDM sufficient even off the path. Reservation payoffs are zero; this too can be realxed and agents can be given the option to opt out of one or more mechanisms.

3. The above observations leads us to ask if there are simple mechanims that are sufficient off the path. A positive answer is provided by our Proposition 3, which shows that a slightly richer class of mechanisms that DDM will suffice even off the path. In additon to reporting the information in a DDM, these sufficient mechanisms report what action the principal being punished is expected to take.

It is well known that repeated games with private monitoring do not admit of parsimonious strategies, even without contracts entering the picture. *A priori* there is no reason to believe that contracts would simplify matters; indeed, in the one-shot model the possibility of writing complicated contracts is responsible for the complexity. However, combining contracts with the dynamic setting pares down the complexity although each model is individually complicated.

Another key difference between the theory of repeated games and our work involves finding minmax values. Computing the minmax in a repeated game is a routine linear programming problem. However in our model the minimum and the maximum must be taken with respect to not principals' actions but contracts mapping into the principals' actions. This class of contracts is intractable. Our equivalence theorems show that the complex minmax of any principal can be computed as his *maxmin* value when he can offer only actions, and the other principals are restricted to *DMs.* so that we can express both the minmax and the equilibrium payoff set in terms of primitives.

3 Setting and Preliminaries

3.1 A Model of Nonexclusive Contracts

The model presented here is an abstraction for repeated many-to-many contracting and trading problems. Theorems are illustrated with numerical examples, with most proofs moved to the appendix.

We first describe the one-shot game that will be repeated over time. The sets of principals and agents[8] are, respectively, $\mathcal{J} := \{1, \cdots, J\}$ and $\mathcal{I} := \{J+1, \cdots, J+I\}$; let $\mathcal{N} = \mathcal{J} \cup \mathcal{I}$. Assume that $I \geq 3$. Each agent i has private information about her (payoff) type θ_i drawn from a finite set Θ_i according to the known distribution μ_i. Let $\theta = (\theta_{J+1}, \dots, \theta_{J+I})$ denote a type profile in $\Theta := \times_{i\in\mathcal{I}}\Theta_i$. Agents' types are independent,[9] i.e. the joint distribution $\mu \in \triangle\Theta$ is the product of the marginals $\mu_i \in \triangle\Theta_i$.[10] Each principal j makes a decision a_j (henceforth referred to as an action) from a finite[11] set A_j. Let $a = (a_1, ...a_J) \in A := \times_{j\in\mathcal{J}} A_j$. A mixed action of principal j is $\alpha_j \in \mathcal{A}_j := \triangle A_j$. Let $\mathcal{A} := \times_{j\in\mathcal{J}}\mathcal{A}_j$, and $\mathcal{A}_{-j} := \times_{k\neq j}\mathcal{A}_k$. For all $n \in \mathcal{N}$, let $u_n : \mathcal{A} \times \Theta \to \mathbb{R}$ be the vNM (von Neumann Morgenstern) expected payoff function for player n, uniformly bounded by $M < \infty$.

The nature and scope of principal j's action a_j are quite general in that it is an allocational decision that principal j and agents can agree on. The interpretation of action and type depends on the application under consideration, as we explain below

- *Multi-unit Trading*: Consider an auction environment where in each period each seller j is endowed with Q units of the good. A buyer can buy mutliple units from multiple sellers. Then, seller j's action $a_j^t = \{(p_{ij}^t, q_{ij}^t)\}_{i\in\mathcal{I}}$ in period t is interpreted as an array of payment and quantity pairs in period t, where q_{ij}^t is the quantity of the good buyer i

[8] We use feminine pronouns for the agent and masculine pronouns for principals.

[9] If types are correlated it makes it easier to extract information about one agent's type from the others as in Crémer and Mclean (1988); see Tripathi (2008) for an application of this idea to dynamic mechanism design with a single principal.

[10] For any set S, the set of probability distributions on S is denoted by ΔS.

[11] Finiteness of the type and action spaces is not critical for our results, but are usually made in the literature. With a modicum of technicalities we can deal with a compact set of actions and a countable type-space.

buys from seller j and p_{ij}^t is buyer i's payment to seller j. Buyer i's type θ_i^t is an array of marginal utilities over successive consumption.

- *Insurance*: Consider risk transfer markets (e.g. credit default swap) where buyers face underlying asset risks over time but their preferences may differ over the desire to shed that risk. For example, buyer i holds an asset that can take one of the two-state contingent values over time: In period t, it returns $r > 0$ with probability $\_\theta_i^t$ (good state) but nothing with probability $1 - \theta_i^t$ (bad state). The quality of the underlying asset in period t is characterized by θ_i^t and it is buyer i's private information. Buyer i can purchase contingent claims that pay in the bad state from multiple sellers. Then each seller j's action $a_j^t = \{(p_{ij}^t, q_{ij}^t)\}_{i\in\mathcal{I}}$ is an array of price and contingent claim pairs in period t, where q_{ij}^t is the amount of contingent claim that buyer i buys from seller j and p_{ij}^t is the price that buyer i pays to seller j for contingent claim q_{ij}^t.

- *Loan Contracting*: Entrepreneur i has a risky investment project that generats profits over time. Let $f(q_i^t)$ be the profit in period t when the amount of money invested in the project is q_i^t and the project turns out to be successful. Let θ_i^t be the probability of success and it is entrepreneus' private information. Entrepreneur i can borrow money from multiple lenders to finance its project. In this case, each lender's action $a_j^t = \{(p_{ij}^t, q_{ij}^t)\}_{i\in\mathcal{I}}$ is an array of pairs of repayment and amount of money borrowed in period t, where p_{ij}^t is the amount of repayment that entrepreneur i makes conditional on the success of her project and q_{ij}^t is the amount of money that entrepreneur i borrows from lender

- *Lobbying*: Several different lobby group act as principals. For example, anti-gun lobby groups and pro-gun lobby groups repeatedly lobby multiple policy makers (e.g., government, Congress, and etc.) for/against gun-control policies over time as issues come along. In this case, each lobby group j's action in period t is denoted by $a_j^t = \{(p_{ij}^t, q_i^t)\}_{i\in\mathcal{I}}$, where $p_{ij}^t$is lobby group j's political support level for policy maker i and q_i^t is the policy decision that policy maker i makes in period t. In this case, q_i^t can be viewed as agent i's contractable task or effort and p_{ij}^t is principal j's reward for it. Policy maker i's ideal policy point in period t can be characterized by θ_i^t and it is her private information.

- *Employment*: An investment advisor often works for multiple investors and an investor may also hire multiple advisors. Employment contracts in period t specifies reward and task for advisors. Investor j's action in period t is denoted by $a_j^t = \{(p_{ij}^t, q_{ij}^t)\}_{i\in\mathcal{I}}$, where q_{ij}^t is advisor i's contractable task that is specific for investor j and p_{ij}^t is investor j's reward for advisor i. A reward can take various form. It could be a function of realized return from investment made by according to the advisor's recommendation or it could be a function of relative ranking of the realized return. Advisor i's expertise or ability in period t can be characterized by her type θ_i^t.

- *Vertical Contracting*: The application in the previous section in fact belongs to vertical contracting where upstream firms (e.g., input producer) contract with downstream firms (e.g., final good producer) Each news agency supplies news reports to multiple news organizations such as newspapers, magazines, and radio and television broadcasters; At the same time, each news organization buys news from multiple new agencies as well. Each cable TV provider contract with multiple TV channels and each TV channel also contract with multiple cable TV providers at the same time. Supply chains in management can be viewed as examples of vertical contracting. Our model can incorporate other examples if they can be modeled as repeated many-to-many trading problems with contracts between multiple principals and multiple agents.

In examples described above, it is seldom observed that a principal (e.g, seller in multi-unit trading problems, seller in insurance problems, lender in loan contracting, lobby group in lobbying, etc.) offers one grand trading mechanism for terms of trade for the entire future, because either principals lack full commitment or all possible future contingencies are too complex to be specified in one grand mechanism. More importantly, when a principal interact with agents repeatedly, it provides him and agents with incentives to repeatedly contract through short-term trading mechanisms for terms of trade over a shorter period of time.Let us formally describe an arbitrary complex mechanism that each principal offers for a contractual decision that applies for a certain period of time. Following Epstein and Peters (1999), a complex mechanism offered by principal j comprises compact message spaces M_{ij} for each $i \in \mathcal{I}$, from which agent i chooses a message for principal j;[12] and a continuous mapping $\gamma_j : M_j \to \mathcal{A}_j$, where $M_j := \times_{i\in\mathcal{I}} M_{ij}$ is the set of all message profiles that might be received by principal j. Each principal's mechanism depends on a potentially very complicated message space, but cannot be directly contingent on the mechanisms of the other principals. The set of all feasible complex mechanisms available to principal j in the stage game is the set of all continuous mappings $\Gamma_j := \{\gamma_j \mid M_j \to \mathcal{A}_j\}$. Let $\Gamma := \times_{j\in\mathcal{J}}\Gamma_j$. Let $M := \times_{j\in\mathcal{J}} M_j$ denote the set of all profiles of messages that might be received by principals collectively.

The timing of the stage game is as follows.

1. The agents observe their respective types $\theta_i \in \Theta_i$.

2. Without observing the agents' types, principals *simultaneously* offer mechanisms γ_j from their respective Γ_j.[13]

[12] The construction of M_{ij} and its compactness are both established in Epstein and Peters (1999).

[13] A principal with full commitment power can offer an infinite-horizon mechanism at the beginning—e.g., Lee and Sabourian (2011); Pavan, Segal and Toikka (2011). We follow Bergemann and Välimäki (2010):

3. After observing the profile of mechanisms offered by principals, each agent sends a private message $m_{ij} \in M_{ij}$ to each principal j. Note that when $m_j \in M_j$ is the profile of messages received by principal j, he chooses the mixed action $\gamma_j(m_j) \in \mathcal{A}_j$, which is observed by agents but not by the other principals.

4. Payoffs are finally earned according to the payoff functions u_n.

We now describe the repeated game $G^{\infty}(\delta)$. We adopt the following notational convention: If $\varkappa$ is a variable in the stage game, we denote its period t value by $\varkappa^t$, with the understanding that t is a superscript and not an exponent. The stage-game is repeated at each time $t \geq 1$. For now we assume that agents' types are independent across periods according to the full support distribution $\mu \in \triangle\Theta$. We significantly relax this assumption later, allowing each agent's type to evolve according to a Markov process.

In keeping with standard mechanism design framework, agents observe the sequence of principals' mechanisms and actions $\{\gamma^s, \alpha^s\}_{s\leq t}$ at the end of period t. To make enforcement harder in our dynamic game, we assume that principals cannot observe mechanisms offered or actions taken by the other principals. However we are able to obtain strong results even in this permissive setting, one weaker than the assumption of perfect monitoring routinely used in repeated games. The discount factor is $\delta \in (0,1)$, i.e. for any $n \in \mathcal{N}$ the (average) discounted payoff of player n from period τ onwards is $(1-\delta)\sum_{t\geq\tau}\delta^{t-\tau}u_n(\alpha^t, \theta^t)$. As usual we assume the existence of a public correlation device (PCD).

3.2 One-shot Mechanisms and Incentive Compatibility

An allocation can be captured by a stage social choice function (stage-SCF), which is a mapping $f\colon \Theta \to \triangle\mathcal{A}$ specifying a probability distribution over action profiles as a function of the type profile of agents. Let $\mathcal{F}$ be the set of all possible stage-SCFs. Given the presence of private information for the agents, it is clear that constant mechanisms, which do not ask for any information from the agents, may not be enough to support f unless f is a constant allocation. The simplest mechanism that elicits the type information is a direct mechanism. We also allow an agent to refuse participation in a mechanism by sending a null message $\emptyset$ to the principal concerned. For each $i \in \mathcal{I}$, let $\tilde{\Theta}_i := \Theta_i \cup \{\emptyset\}$ and $\tilde{\Theta} := \times_i \tilde{\Theta}_i$.

Definition 1 *A direct mechanism (DM) offered by principal j is a mapping $\pi_j\colon \tilde{\Theta} \to \mathcal{A}_j$.*

We provide a few examples where DMs are basic building blocks in commonly observed contracting situations. For example, auctions decide allocations across bidders contingent on bids submitted. Because bids are directly tied to bidders' own willingness to pay, auctions

mechanisms are offered period by period. Since these papers involve a single principal or social planner offering mechanisms, complications introduced by competition among principals does not arise; our work must grapple with this to obtain simple mechanisms.

are essentially DMs. Many incentive contracts can be also characterised by DMs. Let $y_{ij}(\cdot)$ denote principal j's incentive contract that specifies his choice p_{ij} as a function of contractable part of agent i's effort or task q_i. Part of each agent i's effort or task that can be contractable between the agent and each principal j depends on the nature of the agent's effort or task and the scope of contracting in the application. In most seller-buyer problems, q_i is decomposed to $(q_{i1}, \ldots, q_{iJ})$, where q_{ij} is the quantity that buyer i buys from seller j. In this case, an incentive contract $y_{ij}(\cdot)$ is interpreted as a nonlinear price and it often specifies buyer i's payment as only a function of only the quantity that she buys from seller j. In the example of lobbying, lobby group j may want to offer a political contribution scheme $y_{ij}(\cdot)$ to policy maker i to incentivise her policy decision q_i. In this case, a polic decision may not be decomposable so that a political contribution scheme $y_{ij}(\cdot)$ specifies lobby group j political level p_{ij} for policy maker i as a function of policy maker's whole decision q_i. Incentive contracts are also basic contracting buiding blocks in many other economic problems such as loan contracting, insurance, and etc. For any array of incentive contracts $\{y_{ij}(\cdot)\}_{i\in\mathcal{I}}$ that principal j offers, one for each agent, we can find out the equivalent DM $\pi_j = \{\pi^p_{ij}(\cdot), \pi^q_{ij}(\cdot)\}_{i\in\mathcal{I}}$, where $(\pi^p_{ij}(\cdot), \pi^q_{ij}(\cdot))$ determines agent i's payment and quantity as a function of her type report only and, for all $\tilde{\theta}_{ij} \in \tilde{\Theta}_i$, $\pi^p_{ij}(\tilde{\theta}_{ij}) = y_{ij}(\pi^q_{ij}(\tilde{\theta}_{ij}))$.

Let Π_j be the set of all DMs available for principal j in an application under consideration and $\Pi := \times_{j\in\mathcal{J}}\Pi_j$. For example, Π_j includes various formats of auctions for seller j in multi-unit trading problems. In lobbying problems, it includes various political contribution schemes that are legally available for lobby group j. Agent i sends a type report $\tilde{\theta}_{ij}$ to principal j; her profile of type reports is therefore a vector of the form $(\tilde{\theta}_{i1}, \ldots, \tilde{\theta}_{iJ}) \in (\tilde{\Theta}_i)^J$, since there are J principals. Given a profile of DMs $\pi := (\pi_1, \ldots, \pi_J)$, the expected payoff of agent θ_i, when the other agents truthfully report their types, is

$$\mathbb{E}_{\mu_{-i}}\left[u_i(\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \cdots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i}), \theta_i, \theta_{-i})\right],$$

where $\mathbb{E}_{\mu_{-i}}[\cdot]$ is the expectation operator given the probability distribution μ_{-i} over Θ_{-i}.

For any profile of DMs π, the set of all actions profiles induced by truthful type reports is given by $\hat{\mathcal{A}}(\pi) := \{[\pi_1(\theta), \ldots, \pi_J(\theta)] \mid \theta \in \Theta\}$. For any given π we define $B_i(\pi) \subset (\tilde{\Theta}_i)^J$ for each i as the set of all arrays of type reports of agent i, one report to each principal, that lead to an action profile in $\hat{\mathcal{A}}(\pi)$ irrespective of the types of the other agents as long as they report truthfully:

$$B_i(\pi) := \left\{(\tilde{\theta}_{i1}, \ldots \tilde{\theta}_{iJ}) \in (\tilde{\Theta}_i)^J \,\middle|\, \left(\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \ldots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i})\right) \in \hat{\mathcal{A}}(\pi)\ \forall \theta_{-i} \in \Theta_{-i}\right\}.$$

We propose three different definitions of incentive compatibility (IC) below, depending on what kind of misreports are being deterred; this is followed by a comparison of the three notions.

Definition 2 *A profile of DMs $\pi = (\pi_1, \dots, \pi_J)$ satisfies IC over $D_i \subset (\tilde{\Theta}_i)^J$ w.r.t. μ if for all $i \in \mathcal{I}$ and all $\theta = (\theta_i, \theta_{-i}) \in \Theta$ we have*

$$\mathbb{E}_{\mu_{-i}}\left[u_i(\pi(\theta), \theta)\right] \geq \mathbb{E}_{\mu_{-i}}\left[u_i(\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \dots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i}), \theta)\right] \quad \forall\, (\tilde{\theta}_{i1}, \dots, \tilde{\theta}_{iJ}) \in D_i, \qquad (2)$$

where $\pi(\theta) = [\pi_1(\theta_i, \theta_{-i}), \dots, \pi_J(\theta_i, \theta_{-i})]$.

If π satisfies (2) over $D_i = (\tilde{\Theta}_i)^J$ we write $\pi \in \Pi^U$ and say that π is unrestrictedly incentive compatible *(UIC).*

If π satisfies (2) over $D_i = B_i(\pi)$, we write $\pi \in \Pi^C$ and say that π is constrained incentive compatible *(CIC). Clearly $(\theta_i, \dots, \theta_i) \in B_i(\pi)$ for any $\theta_i \in \Theta_i$.*

If π satisfies (2) over $D_i := \{(\tilde{\theta}_{i1}, \dots \tilde{\theta}_{iJ}) \in (\tilde{\Theta}_i)^J \mid \tilde{\theta}_{i1} = \cdots = \tilde{\theta}_{iJ} \in \Theta_i\}$, we write $\pi \in \Pi^W$ and say that π is weak constrained incentive compatible *(WIC).*

For any $K \in \{U, C, W\}$, let Π_j^K be the projection of Π^K onto principal j's space of DMs. For all $\pi_j \in \Pi_j^K$, let $\Pi_{-j}^K(\pi_j) := \{\pi_{-j} \in \Pi_{-j} \mid (\pi_j, \pi_{-j}) \in \Pi^K\}$, i.e. $\Pi_{-j}^K(\pi_j)$ is the set of all $\pi_{-j} \in \Pi_{-j}$ makes (π_j, π_{-j}) KIC given $\pi_j \in \Pi_j^K$. Throughout, IC is defined on the agent's *stage-game* payoff. Given the notion of IC captured in $K \in \{U, C, W\}$, a stage-SCF f is said to be induced by $\pi \in \Pi^K$ if $f(\theta) = \pi(\theta)$ for all $\theta \in \Theta$.[14]

Let us clarify the various notions of IC in terms of the DMs that belong to each class, and the misreports that each notion deters through the use of contemporaneous incentives. Clearly we have $\Pi^U \subset \Pi^C \subset \Pi^W$, because a weaker notion of IC in the above chain needs to deter fewer deviations:

$$(\tilde{\Theta}_i)^J \supset B_i(\pi) \supset \{(\theta_i, \dots, \theta_i) \in (\tilde{\Theta}_i)^J \mid \theta_i \in \Theta_i\} \text{ for any } \pi \in \Pi.$$

Consider an agent i contemplating a unilateral deviation from truthtelling. We say that i reports *consistently* if she reports the same type to all principals ($\tilde{\theta}_{i1} = \cdots = \tilde{\theta}_{iJ}$), and *inconsistently* if $\{\tilde{\theta}_{i1}, \dots, \tilde{\theta}_{iJ}\}$ contains at least two distinct elements, i.e. she sends different messages to at least two principals.

UIC imposes IC for all messages. It is clear that all three notions need IC for all consistent misreports because such lies induce action profiles in $\hat{\mathcal{A}}(\pi)$ and cannot be detected. We have three types of inconsistent reports as follows.

1. Inconsistent messages $(\tilde{\theta}_{i1}, \dots, \tilde{\theta}_{iJ})$ such that $(\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \dots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i})) \notin \hat{\mathcal{A}}(\pi)$ for some θ_{-i} are detected with positive probability. Only UIC imposes IC w.r.t. these.

[14]Note that f is defined on $\times_i \Theta_i$ whereas π is defined on $\times_i \tilde{\Theta}_i$: the latter specifies principals' actions even when some agents do not report their types.

2. Some inconsistent messages $(\tilde{\theta}_{i1}, \dots, \tilde{\theta}_{iJ})$ by i are outcome-equivalent to a consistent lie, i.e. for some $\hat{\theta}_i \in \Theta_i$ we have

$$(\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \dots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i})) = \pi(\hat{\theta}_i, \theta_{-i}) \in \hat{\mathcal{A}}(\pi) \; \forall \, \theta_{-i} \in \Theta_{-i};$$

these create no problems because they are taken care of when we deter consistent lies.

3. The remaining class of messages are the inconsistent reports in $B_i(\pi)$ that lead to an action in $\hat{\mathcal{A}}(\pi)$ but are not equivalent to a single consistent lie. Both UIC and CIC imposes incentive compatibility with respect to these, whereas WIC does not.

Let us explain the different notions of IC with an example below.

Example 1. There are two principals and three agents[15] — $\Theta_1 = \{\theta_1, \theta_1^{'}\}$, $\Theta_2 = \{\theta_2, \theta_2^{'}\}$, $\Theta_3 = \{\theta_3\}$. For simplicity, we assume that agents must participate. Consider a profile of DMs $\pi = (\pi_1, \pi_2)$ that maps type reports to actions as given by the table below. Agent 3's type report is fixed at θ_3; rows correspond to message profiles (one message to each principal) of agent 1, while columns correspond to agent 2's message profiles.

	θ_2, θ_2	$\theta_2, \theta_2^{'}$	$\theta_2^{'}, \theta_2$	$\theta_2^{'}, \theta_2^{'}$
θ_1, θ_1	$\alpha_1, \alpha_2^{'}$	α_1, α_2	$\alpha_1, \alpha_2^{'}$	α_1, α_2
$\theta_1, \theta_1^{'}$	$\alpha_1, \alpha_2^{'}$	$\alpha_1, \alpha_2^{'}$	$\alpha_1, \alpha_2^{'}$	$\alpha_1, \alpha_2^{'}$
$\theta_1^{'}, \theta_1$	$\alpha_1^{'}, \alpha_2^{'}$	$\alpha_1^{'}, \alpha_2$	$\alpha_1^{''}, \alpha_2^{'}$	$\alpha_1^{''}, \alpha_2$
$\theta_1^{'}, \theta_1^{'}$	$\alpha_1^{'}, \alpha_2^{'}$	$\alpha_1^{'}, \alpha_2^{'}$	$\alpha_1^{''}, \alpha_2^{'}$	$\alpha_1^{''}, \alpha_2^{'}$

Truthful reports induce actions profiles in $\hat{\mathcal{A}}(\pi) = \{(\alpha_1, \alpha_2^{'}), (\alpha_1^{'}, \alpha_2^{'}), (\alpha_1, \alpha_2), (\alpha_1^{''}, \alpha_2^{'})\}$. WIC requires that the truthful type report is at least as good as any other consistent type report given the other agents' truthful report. Assume agent 2 reports the truth. If agent 1 reports θ_1 to every principal, she induces $(\alpha_1, \alpha_2^{'})$ and (α_1, α_2) when the other agents' types are (θ_2, θ_3) and $(\theta_2^{'}, \theta_3)$ respectively. However, if agent 1 reports $\theta_1^{'}$ to every principal, it induces $(\alpha_1^{'}, \alpha_2^{'})$ and $(\alpha_1^{''}, \alpha_2^{'})$ when the other agents' types are (θ_2, θ_3) and $(\theta_2^{'}, \theta_3)$ respectively. Under WIC, agent 1's choice is whether to choose $(\alpha_1, \alpha_2^{'})$ and (α_1, α_2), or $(\alpha_1^{'}, \alpha_2^{'})$ and $(\alpha_1^{''}, \alpha_2^{'})$.

Now consider CIC. Given the other agents' truthful type reports, agent 1's inconsistent type report $(\theta_1, \theta_1^{'})$ induces $(\alpha_1, \alpha_2^{'})$ for both possible profiles of the other agents' types. Because $(\alpha_1, \alpha_2^{'}) \in \hat{\mathcal{A}}(\pi)$, the other agents cannot detect agent 1's inconsistent type report $(\theta_1, \theta_1^{'})$ by observing the action profile $(\alpha_1, \alpha_2^{'})$ at the end of the period and hence $(\theta_1, \theta_1^{'}) \in B_1(\pi)$. On the other hand, agent 1's inconsistent type report $(\theta_1^{'}, \theta_1)$ induces $(\alpha_1^{'}, \alpha_2^{'})$ and $(\alpha_1^{''}, \alpha_2)$ when the other agents' types are (θ_2, θ_3) and $(\theta_2^{'}, \theta_3)$ respectively. Because $(\alpha_1^{''}, \alpha_2) \notin$

[15] In this example we refers to agents with $1, 2, 3$ while principals are $1, 2$. The overlap in numbering principals and agents is not problematic in this simple example.

$\hat{\mathcal{A}}(\pi)$, the other agents can detect agent 1's inconsistent type report $(\theta_1^{'}, \theta_1)$ when their type profile is $(\theta_2^{'}, \theta_3)$, i.e., the other agents can detect agent 1's inconsistent type report $(\theta_1^{'}, \theta_1)$ with positive probability. Hence, $(\theta_1^{'}, \theta_1) \notin B_1(\pi)$. In summary, $B_1(\pi) = \{(\theta_1, \theta_1), (\theta_1^{'}, \theta_1), (\theta_1^{'}, \theta_1^{'})\}$.

Given the other agent's truthful type reports, agent 2's inconsistent type report $(\theta_2, \theta_2^{'})$ induces (α_1, α_2) and $(\alpha_1^{'}, \alpha_2^{'})$ when the other agents' types are (θ_1, θ_3) and $(\theta_1^{'}, \theta_3)$ respectively. Because both (α_1, α_2) and $(\alpha_1^{'}, \alpha_2^{'})$ are in $\hat{\mathcal{A}}(\pi)$, agent 2's inconsistent type report $(\theta_2, \theta_2^{'})$ cannot be detected only by observing the action profiles. It means that $(\theta_2, \theta_2^{'}) \in B_2(\pi)$. Similarly, $(\theta_2^{'}, \theta_2) \in B_2(\pi)$ so that $B_2(\pi)$ includes all possible type reports for agent 2: $B_2(\pi) = \{(\theta_2, \theta_2), (\theta_2^{'}, \theta_2), (\theta_2, \theta_2^{'}), (\theta_2^{'}, \theta_2^{'})\}$. CIC is imposed over only the type reports in $B_1(\pi)$ for agent 1 and the type reports in $B_2(\pi)$ for agent 2. While $B_2(\pi)$ includes all possible type reports, $B_1(\pi)$ does not include $(\theta_1^{'}, \theta_1)$ because it can be detected with positive probability. Therefore UIC and CIC are the same for player 2 but not for player 1. ■

In a one-shot model, UIC is the appropriate notion of IC because agents cannot be punished even if a false type report is detected at the end of the game. We will show that in our dynamic setting, IC refers to the weaker notions of CIC and WIC; the legitimacy of these notions rests on our ability to use the continuation game of the repeated game to deter inconsistent messaging. Consider CIC: If agent i's messages does not lie in $B_i(\pi)$, it leads to an action profile outside of $\hat{\mathcal{A}}(\pi)$ with positive probability, given that others report truthfully. Such an action profile informs all agents that at least one of them sent false type reports (although they may not know the identity of this agent). In the dynamic game these lies may be easily deterred because they can be punished with positive probability. On the other hand, it would seem that WIC is too weak because inconsistent messages in $B_i(\pi)$ cannot be identified; we shall return to this towards the end of Section 3.

The next example derives the set Π^K of all profiles of DMs satisfying KIC, depending on the notion of IC captured in $K \in \{U, C, W\}$. This example plays a key role as all further examples build off it.

Example 2. There are two principals and three agents —$\mathcal{A}_1 = \{\alpha_1, \alpha_1^{'}\}$, $\mathcal{A}_2 = \{\alpha_2, \alpha_2^{'}\}$, $\Theta_1 = \{\theta_1, \theta_1^{'}\}$, $\Theta_2 = \{\theta_2\}$, $\Theta_3 = \{\theta_3\}$. Agent 1's type is either θ_1 or θ_1' with equal probability. Agents 2 and 3 have no private information about their types because Θ_2 and Θ_3 are singletons. As in Example 1, assume that agents must participate. Players' payoffs are given by the following tables, one for each possible type of agent 1. The numbers in each cell represent players' payoffs in the following order — principal 1, principal 2, agent 1, agent 2, agent 3.

$\theta = (\theta_1, \theta_2, \theta_3)$

	α_2	$\alpha_2^{'}$
α_1	4,2,2,1,1	3,5,3,1,1,
$\alpha_1^{'}$	6,8,4,1,1	9,9,2,1,1

$\theta^{'} = (\theta_1^{'}, \theta_2, \theta_3)$

	α_2	$\alpha_2^{'}$
α_1	8,6,4,1,1	7,9,2,1,1
$\alpha_1^{'}$	2,3,1,1,1	5,4,3,1,1

For ease of exposition, all examples consider only pure actions and DMs in Π_j map to A_j rather than $\mathcal{A}_j$. For each principal j, there are four possible DMs: $\Pi_j = \{\bar{\pi}_j, \bar{\pi}_j^{'}, \pi_j, \pi_j^{'}\}$, where the four DMs are defined as follows.

$$\bar{\pi}_j(\theta_1, \theta_2, \theta_3) = \alpha_j,\ \bar{\pi}_j(\theta_1^{'}, \theta_2, \theta_3) = \alpha_j$$
$$\bar{\pi}_j^{'}(\theta_1, \theta_2, \theta_3) = \alpha_j^{'},\ \bar{\pi}_j^{'}(\theta_1^{'}, \theta_2, \theta_3) = \alpha_j^{'}$$
$$\pi_j(\theta_1, \theta_2, \theta_3) = \alpha_j,\ \pi_j(\theta_1^{'}, \theta_2, \theta_3) = \alpha_j^{'}$$
$$\pi_j^{'}(\theta_1, \theta_2, \theta_3) = \alpha_j^{'},\ \pi_j^{'}(\theta_1^{'}, \theta_2, \theta_3) = \alpha_j$$

Because agents 2 and 3 have no private information, each mechanism makes the principal's action contingent on agent 1's type report only. Note that $\bar{\pi}_j$ and $\bar{\pi}_j^{'}$ are constant DMs that always assign α_j and α_j' respectively while π_j and π_j' change actions conditional on agent 1's type report. Because each principal has four DMs, there are sixteen profiles of DMs that principals can offer.

Given any profile of mechanisms, agent 1 can report one of four different type profiles in $\Theta_1 \times \Theta_1 = \{(\theta_1, \theta_1), (\theta_1, \theta_1^{'}), (\theta_1^{'}, \theta_1), (\theta_1^{'}, \theta_1^{'})\}$. The notion of UIC imposes incentive compatibility over all possible type reports in $\Theta_1 \times \Theta_1$; it is easy to show that

$$\Pi^U = \{(\bar{\pi}_1, \bar{\pi}_2), (\bar{\pi}_1, \bar{\pi}_2^{'}), (\bar{\pi}_1, \pi_2^{'}), (\bar{\pi}_1^{'}, \bar{\pi}_2), (\bar{\pi}_1^{'}, \bar{\pi}_2^{'}), (\bar{\pi}_1^{'}, \pi_2), (\pi_1, \bar{\pi}_2^{'}), (\pi_1^{'}, \bar{\pi}_2)\}.$$

Now consider profiles of CIC DMs. Incentive compatibility is not needed for a type report that has a positive probability of inducing an action profile that would never happen if all agents were to report truthfully. The table below shows the action profiles induced by agent 1's reports, where rows and columns correspond to the message sent to principals 1 and 2 respectively.

	θ_1	$\theta_1^{'}$
θ_1	$\alpha_1^{'}, \alpha_2$	$\alpha_1^{'}, \alpha_2^{'}$
$\theta_1^{'}$	α_1, α_2	$\alpha_1, \alpha_2^{'}$

Consistent type reports (θ_1, θ_1) and $(\theta_1^{'}, \theta_1^{'})$ induce actions profiles $(\alpha_1^{'}, \alpha_2)$ and $(\alpha_1, \alpha_2^{'})$ respectively and hence $\hat{\mathcal{A}}(\pi_1^{'}, \pi_2) = \{(\alpha_1^{'}, \alpha_2), (\alpha_1, \alpha_2^{'})\}$. The profile of DMs $(\pi_1^{'}, \pi_2)$ is not UIC: If agent 1 of type $\theta_1^{'}$ report $\theta_1^{'}$ to principal 1 but θ_1 to principal 2 so that (α_1, α_2) is induced, then her payoff is 4 while she receives the payoff of 2 from $(\alpha_1, \alpha_2^{'})$ that is assigned by truthful type reports to both principal. However, if we adopt the notion of CIC, we do not need to worry about such an inconsistent type report because (α_1, α_2) is not in $\hat{\mathcal{A}}(\pi_1^{'}, \pi_2)$. In fact, any inconsistent type report results in an action outside $\hat{\mathcal{A}}(\pi_1^{'}, \pi_2)$ with probability one. Hence IC is imposed only over $B_1(\pi_1^{'}, \pi_2) = \{(\theta_1, \theta_1), (\theta_1^{'}, \theta_1^{'})\}$. It is easy to check that $(\pi_1^{'}, \pi_2)$ is CIC; it is also WIC because $B_1(\pi_1^{'}, \pi_2)$ includes only consistent type reports. Similarly, we can show that $(\pi_1^{'}, \pi_2^{'})$ is both CIC and WIC although it is not UIC. Therefore, we have

$$\Pi^C = \Pi^W = \{(\pi_1^{'}, \pi_2), (\pi_1^{'}, \pi_2^{'})\} \cup \Pi^U.$$

Weakening the notion of IC from UIC to CIC or WIC strictly expands the set of profiles of DMs satisfying incentive compatibility. ■

Given a profile of complex mechanisms, each messaging protocol induces a profile of DMs; the next definition formalises this.

Definition 3 *Given any $\gamma = (\gamma_j, \gamma_{-j}) \in \Gamma$, equilibrium or otherwise, a continuation messaging profile $m = (m_1, \ldots, m_J)$, where $m_{ij} : \Theta \to M_{ij}$ and $m_j = (m_{ij})_{i \in \mathcal{I}}$, is said to induce a profile of DMs $\pi = (\pi_1, \ldots, \pi_J)$ if, for all $k \in \mathcal{J}$,*

$$\pi_k(\theta) := \gamma_k\left(m_k\left(\theta\right)\right) \;\; \forall \theta \in \Theta, \tag{3}$$

$$\pi_k\left(\emptyset, \theta_{-i}\right) := \gamma_k\left(\emptyset, m_{-ik}(\theta_{-i})\right) \forall \theta_{-i} \in \Theta_{-i}, \;\; \forall i \in \mathcal{I}. \tag{4}$$

The set of all such DMs is $\Pi(\gamma)$, and $\mathcal{A}\left(\gamma\right)$ is the set of all action profiles that may be induced by γ. Let $\Pi^K(\gamma) \;\; := \Pi(\gamma) \cap \Pi^K$ for any $K \in \{U, C, W\}$, where the superscript K denotes the notion of IC adopted. Finally for any $\gamma_j \in \Gamma_j$ and any $K \in \{U, C, W\}$, define

$$\Pi_j^K(\gamma_j) \;:=\; \bigcup_{\gamma_{-j} \in \Gamma_{-j}} \Pi^K(\gamma_j, \gamma_{-j}).$$

Equation (3) says that the action taken under π_k for truthful reports is the one taken under γ_k if all agents send messages according to m_k; equation (4) says that if agent i refuses participation while the other agents' report valid types, the action taken by π_k is the one taken under γ_k when the agent does not participate in γ_k and the others report according to m_k. The notation above will be useful in defining various minmax values. Note that $\Pi_{-j}^K(\pi_j)$ is a subset of Π_{-j} whereas $\Pi_j^K(\gamma_j)$ is a subset of Π_j.

The dynamic setting involves simpler mechanisms than the static one precisely because mechanisms need not deter contemporaneous deviations. As long as potential deviators can be identified, punishment can be deferred. This motivates the following construction.

Definition 4 *A Deviator-reporting DM (DDM) offered by principal j is a mapping π_j^a : $\times_{i \in \mathcal{I}}[\tilde{\Theta}_i \times \{0, 1 \ldots, J, J+1, \ldots, J+I\}] \to \mathcal{A}_j$.*

For each j, $\pi_j^a(\tilde{\theta}, d_j)$ denotes the action of principal j when the profile of type reports is $\tilde{\theta}$, and the profile of reports on the identity of the deviating player is $d_j = (d_{J+1,j}, \ldots, d_{J+I,j})$. At any time t each agent i reports her type $\tilde{\theta}_{ij}^t \in \tilde{\Theta}_i$, and the identity $d_{ij}^t \in \{0, 1 \ldots, J+1\}$

of the most recent deviator; $d^t_{ij} = 0$ is a report from agent i to principal j that no agent had deviated in the last period $t-1$ and no principals have deviated in the current period t; the report $d^t_{ij} \geq J+1$ amounts to reporting that agent d^t_{ij} had deviated at $t-1$; $0 < d^t_{ij} \leq J$ means that agent i reports that principal d^t_{ij} deviated at t. Note that when agent i sends reports at time t, she is not aware if any agents will deviate in period t; she only knows about principals deviations until t and agents' deviations until $t-1$.

In Yamashita's recommendation mechanisms, each agent's message space is the Cartesian product of all agents' type spaces, which gets exponentially complicated as the number of agents and the cardinality of the type spaces increase. DDMs are much simpler — each agent needs to report only her type and a number between 0 and J. Furthermore Yamashita imposes more demanding incentive compatibility requirements, as we explain later.

The incentive compatibility of DDMs is potentially complex because agents report not only their types but also the deviating player, if any. However, we shall see that incentive compatibility over agents' type reports suffices. However in our equilibrium players reports on types will influence principals' current actions while deviator reports will influence future play; i.e. $\pi^a_j(\tilde{\theta}, d_j) = \pi^a_j(\tilde{\theta}, d'_j) =: \pi_j(\tilde{\theta})$ for all $j \in \mathcal{J}$, $\tilde{\theta} \in \tilde{\Theta}$, $d_j, d'_j \in \{0, 1, \ldots, J+I\}^I$. Thus reporting one's type is KIC under the DDMs $\pi^a = (\pi^a_1, \ldots, \pi^a_J)$ if π satisfies IC. In that sense, agents' reports on the deviator's identity are similar to cheap talk as far as the current period is concerned because the principal's action does not depend on agents' reports on the identity of a deviating player; however, it changes the continuation play.

For example, a DDM characterizes a situation where seller j simply asks bidders if someone deviated in the current period as he offers a multi-unit auction in multi-unit trading problems; Only bids that bidder send will decide seller j's allocation regardless of bidders' reports on who, if any, the deviator is. Also, a DDM characterises a situation where principal j simply asks agents on who, if any, the deviator is as he offers incentive contracts to agents in problems where an incentive contract is a basic contracting building block. In this case, principal j's choices only depend on agents' effort or task but not agents' reports on the identity of the deviator. Because principal$j$'s allocation decision in a DDM depends only on agents' type reports but not on their reports on the identity of the deviator, a DDM is even simpler than price matching practices in which a seller actually has to lower its current price contingent on buyers' reports on lower prices offered by a competing seller.

3.3 Minmax Values

The set of payoffs that can be sustained in equilibrium depends on the minmax values of the players, which in turn depends on the set of mechanisms Γ that principals can offer and the appropriate notion of incentive compatibility. In the general case of interdependent values player $i \in \mathcal{I} \cup \mathcal{J}$ gets utility $u_i(\alpha, \theta)$ when the action profile is α and the type profile is θ. A special case is that of private values, where no agent's type affects any other player's payoff: $u_j(\alpha,\theta) = u_j(\alpha,\theta') \ \forall\alpha \in \mathcal{A}, \forall\theta,\theta' \in \Theta, \forall j \in \mathcal{J}$ and $u_i(\alpha,\theta_i,\theta_{-i}) = u_i(\alpha,\theta_i,\theta'_{-i})$ $\forall\alpha \in A, \forall\theta_i \in \Theta_i, \forall\theta_{-i},\theta'_{-i} \in \Theta_{-i}, \forall i \in \mathcal{I}$. With slight abuse of notation, we use $u_j(\alpha)$ and $u_i(\alpha,\theta_i)$ for principal j's payoff function and agent i's payoff function in the private value case. The mechanisms required to attain the complex minmax values, and hence to replicate the complex mechanism allocations, are simpler under private values.

In the benchmark one-shot game principal j's minmax w_j^1 (relative to Γ) is

$$w_j^1 := \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^1\left(\gamma_{-j},\gamma_j\right), \text{ where } u_j^1(\gamma) := \min_{\pi\in\Pi^1(\gamma)} \mathbb{E}_\mu\left[u_j\left(\pi\left(\theta\right),\theta\right)\right]$$

and $\Pi^1(\gamma)$ is the set of all profiles of UIC DMs that can be induced by all continuation equilibria at $\gamma \in \Gamma$ in the one-shot game. The superscript 1 refers the "one-shot" game. To formulate the set of SCFs that are supported by equilibria of the one-shot game, we let a SCF f be strictly individually rational (SIR) for principals with respect to (w.r.t.) $\mu \in \triangle\Theta$ if

$$\mathbb{E}_\mu\left[u_j\left(f(\theta),\theta\right)\right] > w_j^1 \ \forall j \in \mathcal{J}$$

The set of SCFs supported in equilibria of the one-shot game is given by

$$\mathcal{F}^1(\mu) := \{f \in \mathcal{F} \colon f \text{ is SIR for principals w.r.t. } \mu \text{ and induced by } \pi \in \Pi^U\}. \tag{5}$$

The one-shot game does not explicitly mention an agent's minmax value w_i^1 because agents simply play continuation equilibria of the one-shot game and do not need to be punished; however, the one-shot minmax is implicitly defined as

$$w_i^1 := \inf_{\pi\ \in\mathcal{F}^1(\mu)} \mathbb{E}_\mu\left[u_i\left(\pi,\theta\right)\right] \ \forall i \in \mathcal{I}. \tag{6}$$

Now we formulate players' minmax values in the dynamic game. While CIC and WIC are the right notions of IC in the dynamic game, we formulate minmax values under all notions of IC (UIC, CIC, and WIC) to facilitate comparison with the static minmax value.

Principals' Minmax Values

In order to characterise the set of all feasible equilibrium SFCs relative to any complex Γ, it is important to derive principal j's minmax value w_j^K relative to complex mechanisms Γ, for any given notion of IC captured by the superscript $K \in \{U,C,W\}$. Off the path following principal j's deviation, let $\gamma = (\gamma_{-j},\gamma_j)$ be a profile of complex mechanisms offered by principals. In a continuation equilibrium of γ, agents induce a profile of DMs satisfying KIC. Principal j's minmax is the lowest payoff he receives from among all such DMs: For any

$K \in \{U, C, W\}$, principal j's minmax w_j^K (relative to Γ) is

$$w_j^K := \min_{\gamma_{-j} \in \Gamma_{-j}} \max_{\gamma_j \in \Gamma_j} u_j^K (\gamma_{-j}, \gamma_j), \text{ where } u_j^K(\gamma) := \min_{\pi \in \Pi^K(\gamma)} \mathbb{E}_\mu [u_j (\pi(\theta), \theta)]. \tag{7}$$

The three possible definitions of minmax thus differ in how stringent a notion of IC is imposed on DMs; this relates to which deviations by the agents can be detected and which ones must be contemporaneously deterred through the use of incentives. Equation (7) also clarifies that principal j's minmax depends in general on the set of permissible mechanisms Γ.

The key difference between the static and dynamic settings is that in the latter non-deviating principals use the threat of future punishment to force agents to follow communication protocols violating UIC (but satisfy CIC or WIC); as we will show later, it not only enlarges the set of feasible SCFs but also lowers the minmax values, effectively increasing the severity of punishments for a deviating principal. Since the static game must deter all misreports, it satisfies UIC; it therefore follows that $w_i^1 \geq w_i^U$ for all players $i \in \mathcal{I} \cup \mathcal{J}$. Proposition 1 compares the various minmax values and the payoff sets they can support in equilibrium.

Agents' Minmax Values

What if an agent deviates? In this case, she will be punished by the other players subsequently. In this light, it is important to find the minmax value w_i^K for agent i, given the notion of IC identified in the superscript $K \in \{U, C, W\}$, when principals can offer arbitrary complex mechanisms from Γ; $w_i^K(\theta_i)$ is the minmax when the type of agent i is θ_i. We consider three cases.

Case 1: Reservation Utilities. First, consider the case where agents have reservation payoffs — if agent i does not participate in any mechanism, she earns her reservation payoff $\underline{u}_i(\theta)$ for any $\theta \in \Theta$, independently of principals' actions. We assert that for any $K \in \{U, C, W\}$, $w_i^K(\theta)$ must be equal to $\underline{u}_i(\theta)$. It cannot exceed $\underline{u}_i(\theta)$ as all principals can refuse to have a contractual relationship with the agent, leaving her with $\underline{u}_i(\theta)$; it cannot be lower than $\underline{u}_i(\theta)$ because the agent can refuse participation in all mechanisms. For w_i^1 in the one-shot setting, no agent falls below her reservation payoff under any SCF in $\mathcal{F}^1(\mu)$. Therefore, agent i's complex minmax equals the expected (averaged over all types) reservation payoff for agent i: For any $K \in \{U, C, W\}$,

$$w_i^1 \geq w_i^K = \underline{u}_i := \mathbb{E}_\mu [\underline{u}_i(\theta)] \ \forall i \in \mathcal{I}. \tag{8}$$

Since we allow an agent to send the message $\emptyset$ to each principal and get her reservation payoff, our definition of IC also incorporates a participation constraint.[16]

[16] In the mechanism design literature this is also called individual rationality; we prefer "participation constraint" to distinguish it from "individual rationality" in the sense of repeated games.

Case 2: Interdependent Values without a Reservation Utility. In order to punish agent i we must take into account the types of all agents other than i. For any $K \in \{U, C, W\}$, agent i's minmax value relative to complex mechanisms is agent $i's$ payoff when all principals offer the worst mechanisms and agents other than i coordinate on the worst continuation equilibrium subject to it being incentive compatible (KIC):

$$w_i^K := \min_{\gamma \in \Gamma} \inf_{\pi \in \Pi^K(\gamma)} \mathbb{E}_\mu \left[u_i \left(\pi, \theta \right) \right] \ \forall i \in \mathcal{I}. \tag{9}$$

Note that agent i is indeed best responding in the above definition; the "max" is already incorporated because i's message is a best response to the messages of the others for each DM satisfying KIC.

Case 3: Private Values. For any $K \in \{U, C, W\}$, agent i's minmax value relative to complex mechanisms is given by (9) with θ replaced by θ_i.

When principals minmax agent i with simple actions, she obtains her simple minmax

$$w_i := \mathbb{E}_{\mu_i}[u_i(\alpha^i, \theta_i)], \ \forall i \in \mathcal{I}, \text{ where } \alpha^i \in \arg\min_{\alpha \in \mathcal{A}} \mathbb{E}_{\mu_i} \left[u_i(\alpha, \theta_i) \right]. \tag{10}$$

Lemma 1 (Agents' Minmax) *Without reservation payoffs, the minmax values for each agent i satisfy the following:*

(1) Under interdependent values, for any $K \in \{U, C, W\}$, $w_i^K = \inf_{\pi \in \Pi^K} \mathbb{E}_\mu \left[u_i \left(\pi, \theta \right) \right]$. Furthermore, $w_i^W \leq w_i^C \leq w_i^U \leq w_i^1$ for all $i \in \mathcal{I}$.

(2) Under private values, $w_i = w_i^W = w_i^C = w_i^U \leq w_i^1$.

With reservations payoffs, $\underline{u}_i = w_i^W = w_i^C = w_i^U \leq w_i^1$.

Proof DMs can be constructed below from a profile of complex mechanisms. Let the profile of complex mechanisms $\gamma^i \in \Gamma$ and the messaging strategies (one for each agent h) $m_h^i : \Theta_h \to \times_j M_{hj}$ induce w_i^K in (9). Thus, following γ^i, agents induce actions $\gamma_j^i(m_j^i(\theta))$ by $j \in \mathcal{J}$ when the type profile is θ. Replace the profile γ^i with a profile $\pi^i := (\pi_1^i, ..., \pi_J^i)$ of DMs satisfying

$$\begin{aligned} \pi_j^i \left(\theta \right) &:= \gamma_j^i \left(m_j^i \left(\theta \right) \right) && \forall \theta \in \Theta; \\ \pi_j^i \left(\emptyset, \theta_{-h} \right) &:= \gamma_j^i \left(\emptyset, m_{-hj}^i \left(\theta_{-h} \right) \right) && \forall \theta_{-h} \in \Theta_{-h}, \forall h \in \mathcal{I}. \end{aligned}$$

No agent $h \in \mathcal{I}$ has a profitable deviation under π^i as there was none under γ^i.

Consider each agent's minmax value without reservation payoffs. The following notions of IC are progressively weaker — UIC, CIC and WIC —, and $\Pi^U \supset \mathcal{F}^1(\mu)$, implying that $\Pi^W \supset \Pi^C \supset \Pi^U \supset \mathcal{F}^1(\mu)$. Agent i's minmax values w_i^K and w_i^1 are derived by minimizing her expected payoff over Π^K and $\mathcal{F}^1(\mu)$ respectively; therefore, $w_i^W \leq w_i^C \leq w_i^U \leq w_i^1$.

When there are reservation payoffs, we noted earlier that $\underline{u}_i = w_i^W = w_i^C = w_i^U$. Agent

i's minmax value w_i^1 in the one-shot game cannot be lower than the reservation payoff $\underline{u}_i$; therefore, $\underline{u}_i = w_i^W = w_i^C = w_i^U \leq w_i^1$. ■

Finding simple representations for principals' minmax values is more complicated as they can choose from a rich class of mechanisms. The next two sections study this for interdependent and private values respectively.

4 The General Case: Interdependent Values

We examine the principal's minmax w_j^K defined by (7) given $K \in \{U, C, W\}$. The key question is how to represent this in terms of simple mechanisms. Under interdependent values, the other principals may want agents to reveal their types to them in order to punish principal j more severely. Other principals cannot achieve this by offering DMs instead of complex mechanisms, because the incentive compatibility of the other principals' DMs depends on principal j's mechanism. Principal j, who also wants agents to reveal their types to him when he responds to the other principals, is not *a priori* restricted to simple mechanisms. This could necessitate the use of complex mechanisms by the other principals as well, defeating our goal of simplicity. Fortunately, we shall see that a slight expansion of DMs is enough to replicate each principal's complex minmax value; the principal being punished cannot deviate to a complex mechanism and do strictly better. The various minmax values, will satisfy the ordering $w_j^W \leq w_j^C \leq w_j^U \leq w_j^1$ — as in the case of agents.

Principal's Minmax

The goal of this subsection is to express the principals' minmax values in terms of simpler sets of mechanisms. We begin with the definition of an *action-DM maxmin* for each principal j, which is his maxmin value when he is restricted to mixed actions and all other principals are restricted to DMs that are KIC given the action of j, for $K \in \{U, C, W\}$. The value of introducing this new concept derives from Proposition 1, which shows that principal j's complex minmax w_j^K equals the maxmin value w_j^{K*} for any $K \in \{U, C, W\}$.

Definition 5 (Action-DM Maxmin) *Fix* (G, Θ, μ). *For any* $K \in \{U, C, W\}$*, the* action-DM maxmin *value of* j *is*

$$w_j^{K*} := \max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi_{-j}^K(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta) \right]. \tag{11}$$

For any $j \in \mathcal{J}$, let α_j^{Kj} be the action and $\psi_{-j}^{Kj}(\alpha_j)$ be the corresponding profiles of DMs of the other principals that attains the maxmin value in (11), i.e.

$$\psi_{-j}^{Kj}(\alpha_j) := \{\psi_k^{Kj}(\alpha_j)\}_{k \neq j} \in \operatorname*{arg\,min}_{\pi_{-j} \in \Pi_{-j}^K(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta) \right],$$

where the superscript denotes the player whose maxmin we are considering.

Example 3 (Continuation of Example 2). Our goal is to compute principal 1's pure minmax value with respect to complex mechanisms, via his action-DM pure maxmin value. When principal 1 deviates, the only non-deviating principal is 2; therefore, the DMs of principal 2 that are KIC conditional on principal 1's action do not vary with K. Fixing principal 1's action at α_1, the set of DMs of principal 2 that are KIC conditional on α_1 is $\Pi^K_{-1}(\alpha_1) = \{\bar{\pi}_2, \bar{\pi}'_2, \pi'_2\}$ for all $K \in \{U, C, W\}$. Note that π_2 is not IC given α_1: If principals offer (α_1, π_2) and agents 2 and 3 report truthfully, agent 1 of type θ_1 can receive the payoff 3 by reporting θ'_1 to principal 2, which is strictly greater than the payoff of 2 from reporting her type truthfully. Given α_1, DMs in $\Pi^K_{-1}(\alpha_1)$ generate expected payoffs for principal 1 as follows:

$$\mathbb{E}_\mu \left[u_1(\bar{\pi}_2(\theta), \alpha_1, \theta)\right] = \frac{1}{2} \times 4 + \frac{1}{2} \times 8 = 6,$$
$$\mathbb{E}_\mu \left[u_1(\bar{\pi}'_2(\theta), \alpha_1, \theta)\right] = \frac{1}{2} \times 3 + \frac{1}{2} \times 7 = 5,$$
$$\mathbb{E}_\mu \left[u_1(\pi'_2(\theta), \alpha_1, \theta)\right] = \frac{1}{2} \times 3 + \frac{1}{2} \times 8 = 5.5.$$

Therefore, if principal 1 plays α_1, the DM $\bar{\pi}'_2$ minimises principal 1's expected payoff from among all DMs in $\Pi^K_{-1}(\alpha_1)$:

$$\psi^{K1}_2(\alpha_1) = \bar{\pi}'_2, \text{ and } u_1(\alpha_1, \bar{\pi}'_2) = 5.$$

Similarly, $\Pi^K_{-1}(\alpha'_1) = \{\bar{\pi}_2, \bar{\pi}'_2, \pi_2\}$. Since $\bar{\pi}_2$ minimises principal 1's expected payoff among all DMs in $\Pi^K_{-1}(\alpha'_1)$, we have

$$\psi^{K1}_2(\alpha'_1) = \bar{\pi}_2, \text{ and } u_1(\alpha'_1, \bar{\pi}_2) = \frac{1}{2} \times 6 + \frac{1}{2} \times 2 = 4.$$

Since $u_1(\alpha'_1, \bar{\pi}_2) = 4 < u_1(\alpha_1, \bar{\pi}'_2) = 5$, it follows that

$$\alpha^{K1}_1 = \alpha_1, \text{ and } w^{K*}_1 = 5 \text{ for all } K \in \{U, C, W\}.$$

The main implication of Proposition 1 is that agent 1's pure minmax value with respect to complex mechanisms is the same as his pure action-DM maxmin action $w^K_1 = w^{K*}_1 = 5$. No such algorithm exists in the one-shot model.

Similarly, for principal 2: $\alpha^{K2}_2 = \alpha'_2$, and $w^{K*}_2 = 4.5$ for all $K \in \{U, C, W\}$. ■

Lemma 2 shows that for any profile γ_{-j} offered by principals other than j in any equilibrium where agents pick the worst KIC continuation for j, one of the best responses of principal j is a constant mechanism. Note that this is true only if agents pick the worst current continuation equilibrium; if agents were to best-respond in the static game, principal

j could improve his payoff by offering a more complicated contract. Agents incentives to play in this manner (rather than play a myopic best-response) derives incentives in the dynamic game and are addressed in the folk theorems.

Lemma 2 *For any $K \in \{U, C, W\}$ and $\gamma_{-j} \in \Gamma_{-j}$,*

$$\max_{\gamma_j \in \Gamma_j} u_j^K(\gamma_{-j}, \gamma_j) = \max_{\alpha_j \in \mathcal{A}_j} u_j^K(\gamma_{-j}, \alpha_j).$$

Proof See the appendix. ■

In the above construction agents induce the action

$$g^j(\gamma_j) := \arg\min_{\alpha_j \in \mathcal{A}_j(\gamma_j)} \left\{ \min_{\pi_{-j} \in \Pi_{-j}^K(\gamma_{-j}, \alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta) \right] \right\} \tag{12}$$

from any γ_j that j might offer. Agents incentives to do so will be derived in the folk theorems.

Off the path, an agent's incentive to report her true type depends on principal j's action; it is therefore not enough for the other principals to offer DMs instead of complex mechanisms. Off the path following principal j's deviation, each principal $k \neq j$ must ask an agent to report three units of information — (i) her type; (ii) the action that principal j will play in the current period if all agents were to follow a prescribed messaging strategy; (iii) the identity of the most recent deviator. Intuitively, our construction provides a way for the other principals to neutralise principal j's ability to condition his action on agents' messages; this is done by forcing agents to induce an action from j independently of their types, but depending on the mechanism principal j offers. This makes one of principal j's best responses an element of $\mathcal{A}_j$.

Definition 6 *An extended direct mechanism (EDM) offered by k when j is being punished is denoted by*

$$\lambda_k^j : E_{kj} \to \mathcal{A}_k \text{ with } E_{ikj} = \mathcal{A}_j \times \{0, 1, \ldots J + I\} \times \tilde{\Theta}_i \text{ and } E_{kj} = \times_i E_{ikj}.$$

The typical element of E_{ikj} is $e_{ikj}^t = (p_{ikj}^t, d_{ik}^t, \tilde{\theta}_{ik}^t)$, where $p_{ikj}^t \in \mathcal{A}_j$, $d_{ik}^t \in \{0, 1, \ldots, I + J\}$, $\tilde{\theta}_{ik}^t \in \tilde{\Theta}_i$. EDMs, which are used only if a principal deviates, include punishments. In contrast, DDMs are offered on the path; these do not include punishments but ask agents to report the identity of the deviating player.

For notational consistency, a caret is used here to denote the majority value of a variable; thus, for example, $\hat{p}_{kj}^t$ is the majority value from $p_{kj}^t = (p_{ikj}^t : i \in \mathcal{I})$, and so on. Let $\tilde{\theta}_k = (\tilde{\theta}_{J+1,k}, \ldots, \tilde{\theta}_{J+I,k})$ and $d_k^t = (d_{J+1,k}^t, \ldots, d_{J+I,k}^t)$. For any $K \in \{U, C, W\}$, we define

$$\lambda_k^{Kj}(p_{kj}^t, d_k^t, \tilde{\theta}_k) := \psi_k^{Kj}(\hat{p}_k^t)(\tilde{\theta}_k) \; \forall \, (p_{kj}^t, d_k^t, \tilde{\theta}_k) \in E_{kj}. \tag{13}$$

Example 4 (Continuation of Examples 2 and 3). We construct principal 1's EDM λ_1^{K2} following principal 2's deviation. Each agent reports principal 2's action, the identity of the most recent deviator, and her type. If two or more agents report α_2, then principal 1's EDM determines his action conditional on agents' type reports according to the DM $\psi_1^{K2}(\alpha_2) = \bar{\pi}_1$. If two or more agents report $\alpha_2^{'}$, then principal 1's EDM determines his action conditional on agents' type reports according to the DM $\psi_1^{K2}(\alpha_2^{'}) = \pi_1$. Note that agents' reports on the identity of the most recent deviator do not affect principal 1's action decision in his EDM. We can similarly construct principal 2's EDM off the path following principal 1's deviation.■

Given EDMs with any $K \in \{U, C, W\}$, Proposition 1 provides a simple algorithm to calculate the minmax value with respect to the Epstein-Peters mechanisms Γ, without having to compute the message spaces M or the mappings in Γ.

Proposition 1 (Minmax Equivalence) *For any $K \in \{U, C, W\}$ and any principal j, $w_j^K = w_j^{K*}$, as specified in (11). It is attained in the truth-telling continuation equilibrium where the other principals offer EDMs λ_{-j}^{Kj} and principal j offers the constant mechanism playing α_j^{Kj}.*

Proof See the appendix. ■

Since applications sometimes restrict attention to the analytically simpler case with all mechanisms (a degenerate case of which is a constant action) mapping from messages to pure actions, we provide a pure-action counterpart of the earlier result. The proof of Proposition 1 applies unmodified save for mixed actions being replaced by pure actions.

Corollary 1 (Pure Action Minmax Equivalence) *If equation (11) is redefined so that the* max *and the* min *are taken over, respectively, pure actions and mappings from messages to pure actions, the equivalence result continues to apply: $w_j^* = w_j^{K*}$.*

Finally, we compare the principal's minmax values across different notions of IC.

Proposition 2 *For any principal j, $w_j^W \leq w_j^C \leq w_j^U \leq w_j^1$.*

Proof See the appendix. ■

Note that principal j's action in a continuation equilibrium off the path is characterised by a DM $\pi_j \in \Pi_j^K$ that embodies a particular notion of IC. In the dynamic game, principal j cannot make his action contingent on agents' types in the worst continuation equilibrium in which agents punish principal j most severely. Hence, only the fixed actions or constant DMs are feasible for principal j from Π_j^K for all $K \in \{U, C, W\}$. In the one-shot game, agents still play the worst continuation equilibrium, and hence some UIC DMs in Π_j^U may never be played. However, we do not know *a priori* which ones are excluded, and therefore the value of the minmax w_j^1 in the one-shot game of Yamashita (2010) is not expressible in a simpler form of DMs.

However, we were still able to show that w_j^1 is higher than w_j^K for all $K \in \{U, C, W\}$. Intuitively, there are two reasons. First, any continuation equilibrium off the path in Yamashita's one-shot game requires UIC for all principals including the deviating principal and hence non-deviating principals cannot force agents to induce the same action from the deviating principal's mechanism regardless of their types: Since this is possible in our dynamic model, the principal being punished cannot do any better by offering a complex mechanism instead of a simple action. Therefore, even with the same notion of UIC, we have $w_j^U \leq w_j^1$.

Second, KIC for all $K \in \{U, C, W\}$ is required for non-deviating principals only given the deviating principals' action in our dynamic model rather than for all principals. In a model with two principals, only one non-deviating principal offers a mechanism to agents in order to punish a deviating principal. As we showed in Example 3, when there are only two principals all notions of IC are identical for the non-deviating principal's DM, i.e., $\Pi_{-j}^W(\alpha_j) = \Pi_{-j}^C(\alpha_j) = \Pi_{-j}^U(\alpha_j)$ for all α_j; therefore, $w_j^W = w_j^C = w_j^U$ in models with two principals.[17] However, with three or more principals the various notions of IC are generally distinct present for a profile of DMs for multiple non-deviating principals. Because the notion of IC is progressively weaker in the order of UIC, CIC and WIC, we can generally establish $w_j^W \leq w_j^C \leq w_j^U$.

4.1 Folk Theorems for the General Case

4.1.1 I.I.D. Types

Now we are ready to establish the folk theorem for the case of interdependent values. What stage-SCFs and payoff profiles can we support in a perfect Bayesian equilibrium (PBE) of $G^\infty(\delta)$ relative to Γ? Consider i.i.d. types, where $\mu \in \triangle\Theta$ is the product of independent distributions. For each $K \in \{U, C, W\}$, the stage-SCF f is SIR w.r.t. μ if each player gets an expected payoff above his/her minmax:

$$\mathbb{E}_\mu \left[u_n \left(f(\theta), \theta\right)\right] > w_n^K \text{ for all } n \in \mathcal{I} \cup \mathcal{J}. \tag{14}$$

For $n \in \mathcal{J}$ and $K \in \{U, C, W\}$, the minmax value w_n^K equals w_n^{K*}; for $n \in \mathcal{I}$, it is equal to (8) or (9), depending on the situation. The SCF f is weakly individually rational (WIR) for i with respect to (w.r.t.) $\mu \in \triangle\Theta$ if $\mathbb{E}_\mu \left[u_i \left(f(\theta), \theta\right)\right] \geq w_i^U$. Define

$$\mathcal{F}^K(\mu) := \{f \in \mathcal{F} \mid f \text{ is SIR w.r.t. } \mu \text{ and induced by } \pi \in \Pi^K\}.$$

The example below shows how to compute $\mathcal{F}^K(\mu)$.

[17] Even with two principals, the inequality $w_j^U \leq w_j^1$ cannot in general be written as an equality for reasons described in the previous paragraph.

Example 5 (Continuation of Example 2). Because agents must participate in all mechanisms, a profile of DMs is also an SCF. Eight profiles of DMs satisfy UIC and therefore lie in Π^U; CIC and WIC admit two additional profiles of DMs. The table below shows expected payoffs for principals 1 and 2 at each profile in $\Pi^C = \Pi^W$ when agents report truthfully; the first eight profiles of DMs are UIC (and therefore CIC and WIC) but the last two profiles are CIC (and WIC) but not UIC.

	$\bar{\pi}_1, \bar{\pi}_2$	$\bar{\pi}_1, \bar{\pi}_2^{'}$	$\bar{\pi}_1, \pi_2^{'}$	$\bar{\pi}_1^{'}, \bar{\pi}_2$	$\bar{\pi}_1^{'}, \bar{\pi}_2^{'}$	$\bar{\pi}_1^{'}, \pi_2$	$\pi_1, \bar{\pi}_2^{'}$	$\pi_1^{'}, \bar{\pi}_2$	$\pi_1^{'}, \pi_2$	$\pi_1^{'}, \pi_2^{'}$
P1	6	5	5.5	4	7	7.5	6	7	6.5	8.5
P2	3	7	5.5	5.5	6.5	6	4.5	7	8.5	7.5

Principals' minmax values are $w_1^K = w_1^{K*} = 5$ and $w_2^K = w_2^{K*} = 4.5$ for all $K \in \{U, C, W\}$. Because agent must participate, we need to consider strict individual rationality for principals only; therefore,

$$\mathcal{F}^U(\mu) = \{(\bar{\pi}_1, \pi_2^{'}), (\bar{\pi}_1^{'}, \bar{\pi}_2^{'}), (\bar{\pi}_1^{'}, \pi_2), (\pi_1^{'}, \bar{\pi}_2)\},$$
$$\mathcal{F}^C(\mu) = \mathcal{F}^W(\mu) = \{(\bar{\pi}_1, \pi_2^{'}), (\bar{\pi}_1^{'}, \bar{\pi}_2^{'}), (\bar{\pi}_1^{'}, \pi_2), (\pi_1^{'}, \bar{\pi}_2), (\pi_1^{'}, \pi_2), (\pi_1^{'}, \pi_2^{'})\}.$$

Clearly, $\mathcal{F}^C(\mu)$ and $\mathcal{F}^W(\mu)$ include two SFCs, $(\pi_1^{'}, \pi_2)$ and $(\pi_1^{'}, \pi_2^{'})$ not included in $\mathcal{F}^U(\mu)$. In particular, both principals strictly prefers $(\pi_1^{'}, \pi_2^{'})$ to any other SCFs in $\mathcal{F}^U(\mu)$. ■

We first present the folk theorem using CIC, and use it to derive an approximate folk theorem for WIC. The theorem below shows that any SCF $f \in \mathcal{F}^C(\mu)$ is supportable in a PBE of $G^\infty(\delta)$ relative to Γ, provided players are sufficiently patient. Principals offer a profile of DDMs $\pi^a = (\pi_1^a, \ldots, \pi_J^a)$ such that

$$\pi_k^a(\theta, d_k) := f_k(\theta) \quad \forall\, d_k \in \{0, 1, \ldots, J+I\}^J,\ \forall\, \theta \in \Theta.$$

If principals continue offering π^a, play a truthful continuation equilibrium in which agents report their true types and 0 as the identity of the deviator. If principal j unilaterally deviates from the agreement at time t, agents report j and play a (previously agreed on but otherwise arbitrary) continuation equilibrium of (γ_j, π_{-j}^a) in the current period; at $t+1$ the other principals offers the EDMs defined in (13) to minmax j. Off the path following principal j's deviation, Proposition 1 shows that principal j can best respond with an action if agents follow the above protocol.

We make the standard full dimensionality assumption **(FD)** that set of expected payoffs is full dimensional, i.e. $\dim[u(\mathcal{F}^C(\mu))] = J + I$. **Caveat:** In discussions related to folk theorems, generic players are denoted by i and j unless explicitly noted otherwise.

Definition 7 *Fix $f \in \mathcal{F}^C(\mu)$. A family of vectors $\{\beta^1, \ldots, \beta^{J+I}\} \subset S^C$ is said to be a*

PSP (Player-Specific Punishment) for the target payoff $v = u(f)$ if it satisfies the following properties $\forall i, j \in \mathcal{I} \cup \mathcal{J}$:

1. *strict individual rationality (SIR): $\beta^i_j > w_j$;*
2. *target payoff domination: $\beta^i_j < v_j$;*
3. *payoff asymmetry (PA): $\beta^i_i < \beta^j_i$ if $i \neq j$.*

Lemma 3 *Fix $f \in \mathcal{F}^C(\mu)$. There exists a family of $I + J$ profiles of mechanisms $\{\pi^{ai} : \tilde{\Theta} \to \mathcal{A} \mid i \in \mathcal{N}\}$ such that each π^{ai}_k is one-to-one and the family $\{\beta^i := \mathbb{E}_\mu u\left(\pi^{ai}(\theta)\right) \mid i \in \mathcal{N}\}$ is a PSP for $v = u(f)$.*

Proof See the appendix. ∎

Proposition 3 (Exact Folk Theorem for i.i.d. Types) *Consider i.i.d. types with distribution $\mu \in \triangle\Theta$. Under FD in a model with interdependent values, any SCF $f \in \mathcal{F}^C(\mu)$ is the outcome of a PBE of $G^\infty(\delta)$ relative to any Γ for high δ. DDMs suffice on path, while off the path following a deviation by principal j, all other principals employ EDMs while j offers a constant mechanism.*

Proof See the appendix. ∎

Corollary 2 *If $f \in \mathcal{F}^W(\mu)$ is induced by an invertible profile $\pi \in \Pi$, the above result applies since all inconsistent reports are detected with certainty.*

The next proposition weakens the notion of IC from CIC to WIC.[18] In this case we are able to virtually support a social choice function in the following sense.[19] Let us write down any given SCF f as $f = (f_1, \ldots, f_J)$, where f_j is a mapping from Θ to $\mathcal{A}_j$. Consider two SCFs f and $\hat{f}$. Fix a type profile to θ. Given two distributions $f_j(\theta)$ and $\hat{f}_j(\theta)$ on A_j, define the distance between the two as

$$\| f_j(\theta) - \hat{f}_j(\theta) \| := \sum_{a_j \in A_j} | f_j(\theta)(a_j) - \hat{f}_j(\theta)(a_j) |;$$

Then, the norm of f and $\hat{f}$ is defined as $\| f - \hat{f} \| := \sum_{\theta \in \Theta} \sum_{j \in \mathcal{J}} | f_j(\theta) - \hat{f}_j(\theta) |$.

Definition 8 *An SCF f is said to be virtually supported if for any $\varepsilon > 0$ we can find an SCF f^ε that can be supported in a PBE of $G^\infty(\delta)$ and satisfies $\| f^\varepsilon - f \| < \varepsilon$.*

[18] While both CIC and WIC are identical in Example 2, in general there will be SCFs that are WIC but not CIC.

[19] Our reasons for virtually, rather than exactly, supporting f are very different from those for virtual implementation in Matsushima (1988) and Abreu and Matsushima (1992).

Proposition 4 (Approximate Folk Theorem for i.i.d. Types) *Let types be i.i.d. $\mu \in \triangle\Theta$. Let FD hold for $u(\mathcal{F}^W(\mu))$. Under interdependent values, for any f satisfying WIC there exists $\underline{\delta} \in (0,1)$ such that, for any $\delta \geq \underline{\delta}$, f can be virtually supported in a PBE of $G^\infty(\delta)$ relative to any Γ using only DDMs on the path and EDMs off the path.*

Proof See the appendix. ■

Let us provide some intuition for the proof. Suppose we take an SCF f satisfying all the constraints for WIC with strict inequality. Then we can perturb each f_j in such a way that the IC conditions are not altered but the perturbed version is invertible and hence satisfies CIC. Then we can apply the previous result. However this procedure fails whenever some of the IC conditions are satisfied with equality, because the perturbation might go the wrong way and violate an IC condition. To get around this we adopt a two-step procedure. The first step makes use of an observation from Abreu and Matsushima — we can find lotteries over actions in A, one lottery for each type of each agent i, in such a way that if each any player i is told that if she reports her type as θ_i then the lottery chosen will be the one that is optimal for i from among all the lotteries on offer. Now modify f with a small probability that these lotteries would be played. This is a perturbation which makes all IC constraints hold with strict inequality. The second step is to perturb this SCF so as to make it invertible without reversing any of the IC inequalities.

4.1.2 Comparison with Other Competing Mechanism Games

In the one-shot setting in Yamashita (2010), the minmax of principal j is given by w_j^1 under the notion of UIC, the most stringent notion of IC, because each profile of DMs arises from a continuation equilibrium $m_k = (m_{1k}, \dots, m_{Ik})$ of the one-shot game and must thus lie in $\Pi^1(\gamma)$. A simple characterisation of w_j^1, and therefore of $\mathcal{F}^1(\mu)$, is not available. Clearly, $\mathcal{F}^1(\mu) = \mathcal{F}_+^1(\mu) \cup \mathcal{F}_=^1(\mu)$, where at least one agent gets her minmax value in $\mathcal{F}_=^U(\mu)$, while everyone gets strictly above his/her minmax value in the other set.

Proposition 5 $\mathcal{F}_+^1(\mu) \subset \mathcal{F}^U(\mu) \subset \mathcal{F}^C(\mu) \subset \mathcal{F}^W(\mu)$.

Proof Agent i's minmax values satisfy $w_i^W \leq w_i^C \leq w_i^U$ with/without reservation payoffs (Lemma 1) and principal j's minmax values satisfy $w_j^W \leq w_j^C \leq w_j^U$ according to Proposition 2. Secondly, $\Pi^W \supset \Pi^C \supset \Pi^U$ because WIC is weaker than CIC, which is weaker than UIC. This leads to $\mathcal{F}^U(\mu) \subset \mathcal{F}^C(\mu) \subset \mathcal{F}^W(\mu)$. Similarly, $\mathcal{F}_+^1(\mu) \subset \mathcal{F}^U(\mu)$ follows from $w_j^U \leq w_j^1$.
■

Any SCF in $\mathcal{F}^C(\mu)$ is supportable in equilibria of the dynamic game with sufficiently patient players. The static game can support any $f \in \mathcal{F}^1(\mu)$. To show that any f that is supported in the static game is also supported in the dynamic game we need to show that any $f \in \mathcal{F}^1_=(\mu)$ can be supported in the dynamic game with sufficiently patient players. The PSP constructed in Lemma 3 do not work in this setting because if an agent is already at her minmax value w_i^W we cannot lower it further.

Proposition 6 (Static versus Dynamic Game) *If $f \in \mathcal{F}^1(\mu)$ then f can be supported in a PBE of the repeated game with low discounting using only DDMs on path and EDMs off path.*

Proof If $f \in \mathcal{F}^1_+(\mu)$ the result follows immediately from Propositions 3 and 5. If $f \in \mathcal{F}^1_=(\mu)$, let $v = u(f)$, and define $\mathcal{I}_= := \{i \in \mathcal{I} \mid v_i = w_i^1\}$ as the set of all agents who earn exactly their minmax values. Find a feasible point v^* such that $v_j^* < v_j$ for all $j \in \mathcal{J}$, $v_i^* = v_i$ for all $i \in \mathcal{I} \setminus \mathcal{I}_=$, and $v_i^* > v_i$ for all $i \in \mathcal{I}_=$. Since v^* is SIR, use Proposition 3 to find $\delta_{\min}$ such that if $\delta \geq \delta_{\min}$ there is an equilibrium of $G^\infty(\delta)$ with payoff v^*. Ignore any deviations by agents if no principals have deviated before. If any principal deviates, then play the strategies that give payoff v^*. In other words we move to a SIR continuation payoff point that punishes all principals, while rewarding those agents who were previously held at their minmax levels. Clearly, no principal has an incentive to deviate from the strategies that give a payoff of v if $(1-\delta)\bar{M} + \delta v_j^* \leq v_j$. The proposition is valid for $\delta \geq \max\{\delta_{\min},\ (\bar{M} - v_j)/(\bar{M} - v_j^*)\}$. ■

The dynamic setting is able to support more social choice functions for two reasons. First, a weaker notion of IC admits additional profiles of DMs. Second, relaxing the notion of IC lowers players' minmax values.

Finally, there exists an equivalence between two models — one where principals simultaneously offers DMs $\pi = (\pi_1, \ldots, \pi_J)$ satisfying WIC and the other where a single grand mechanism $\pi^G : \tilde{\Theta} \to \mathcal{A}$ is offered — in terms of the set of social choice functions that can be supported. Instead of each principal j offering his own DM $\pi_j \colon \tilde{\Theta} \to \mathcal{A}_j$, suppose that a third party offers only one grand DM $\pi^G : \tilde{\Theta} \to \mathcal{A}$ and each agent sends a single type report to the third party only. It is clear that an equilibrium profile of DMs π must at least be WIC. By extracting all consistent type reports from π, we can construct the corresponding single grand mechanism π^G where truthful type reports support the same SCF as in π under truthful reports.

4.1.3 Markov Types

Propositions 3 and 4 hold when types are i.i.d. draws from the measure space $(\Theta, 2^\Theta, \mu)$. A more general assumption would be that the types form a Markov chain with the transition matrix P and an initial distribution $\mu^0 \in \triangle\Theta$, where

$$P = [p_{rs}]_{r,s \in \Theta}\,, \quad \text{with } \sum_s p_{rs} = 1\ \forall r,\ p_{rs} \geq 0.$$

The key step in extending our results to this setting involves specifying the payoff of an agent along any path, either the original equilibrium path or a punishment path. We start with the following theorems, well known from the literature on stochastic processes (see Breiman (1991) for definitions and proofs.)

Theorem 1 *A finite-state Markov chain $(X)_{t\geq 0}$ with a transition matrix $P = [p_{rs}]_{r,s\in\Theta}$ is irreducible iff it has a unique stationary distribution*[20] *μ^* such that $\mu^* P = \mu^*$.*

A sufficient condition for this is a full-support assumption often employed in game theoretic models: $p_{rs} > 0\ \forall r, s$, i.e. all entries of the transition matrix are strictly positive.

Theorem 2 *If a Markov chain $(X)_{t\geq 1}$ with a transition matrix P taking values in the finite state space $(\Theta, 2^{\Theta})$ possesses a unique stationary distribution μ^* then*

$$\mathbb{P}(X^t = \theta) \to \mu^*(\theta) \ \textit{as}\ t \to \infty\ ; \theta \in \Theta$$

iff the Markov chain is aperiodic.

First of all, let us state what IC means in this setting. When types follow a Markov process and the period-t type profile for agents except for i is θ_{-i}, the distribution over the type profiles in Θ_{-i} at time $t+1$ is given by $\mu_{-i}(\acute{\theta}_{-i} \mid \theta_{-i}) := p_{\theta_{-i}\acute{\theta}_{-i}}$. We use the matrix P_{-i} to represent the transition matrix for players other than i; in other words, each row is a probability distributions over the next period's types conditional on a current type profile in Θ_{-i}. Let $\mathcal{R}(P_{-i})$ denote the row space of this matrix. Given any beliefs i has about the current types of other agents, his belief over the next period's type vector for the others lies in $\mathcal{R}(P_{-i})$. Let the notion of IC be captured by the superscript $K \in \{MU, MC, MW\}$, denoting Markov UIC, Markov CIC and Markov WIC respectively.

Since the distribution over types is changing over time, we are naturally led to ask: With respect to which distribution should we take expectations in (14)? It turns out that the right distribution is the steady-state distribution μ^* from Theorem 1. If types follow independent Markov chains with the steady state distribution μ^* we have to define the minmax values accordingly. For principal j, we define w_j^K for each $j \in \mathcal{J}$ by first deriving a minmax value $w_j^K(p^l)$ for each row p^l of the transition matrix for $l \in \{1, \ldots, L\}$. Now we define $w_j^K := \sum \mu_l^* w_j^K(p^l)$. For agent i, we can define w_i^K similar to w_j^K without reservation payoffs. When there are reservation payoffs, w_i^K is based on equation (8) with the steady-steady state distribution μ^*. Now let the stage-SCF f be SIR w.r.t μ^* with the proper minmax values given a notion of IC.

[20] For a Markov chain with a *countable* state-space, we also need *positive recurrence* to guarantee the existence of a stationary distribution. If the set of types is *finite*, as in most applied work, irreducibility implies positive recurrence.

Definition 9 *Let $K \in \{MU, MC, MW\}$ be the notion of IC adopted in the Markov type case. A profile of DMs π is said to satisfy Markov KIC if the respective notion of IC holds for each i and each conditional distribution $\mu_{-i}(\cdot \mid \theta_{-i}) \in P_{-i}$. Let Π^{MK} be the set of all profiles of Markov incentive compatible (KIC) DMs. Define*

$$\mathcal{F}^{MK}(\mu^*) := \{f \in \mathcal{F} \mid f \text{ is SIR w.r.t. } \mu^* \text{ and induced by } \pi \in \Pi^{MK}\}.$$

Before we explain why our definition of Markov CIC (and Markov WIC) suffices, we state the main theorem of this subsection with the notion of Markov CIC.

Proposition 7 (Folk Theorem with Markov Types) *Let types evolve according to independent irreducible Markov chains in a model of interdependent values where $u(\mathcal{F}^{MC}(\mu^*))$ satisfies FD. Any SCF $f \in \mathcal{F}^{MC}(\mu^*)$ is the outcome of a PBE of $G^\infty(\delta)$ relative to any Γ for high δ; in these PBE, it suffices to use DDMs on path, while off the path following a deviation by principal j, all other principals can use EDMs while j offers a constant mechanism.*

Proof See the appendix. ∎

Remark Similar to Proposition **4**, we can weaken CIC to virtually support an SCF that satisfies Markov WIC but not Markov CIC.

In the Markov case, agent i's IC should be based on her probabilistic belief about the current types of other agents. If θ_{-i}^{t-1} were known, i would use the distribution $\mu_{-i}(\cdot \mid \theta_{-i}^{t-1}) \in \triangle\Theta_{-i}$ to take expectations. But agent i may not know θ_{-i}^{t-1} and assigns a probability to each row of the transition matrix P_{-i}. Her belief over types in Θ_{-i} is thus a convex combination of rows of P_{-i}, i.e. an element of $\mathcal{R}(P_{-i})$. It is easy to check that the set of probability distributions over which IC hold is a convex set. Thus our definition of Markov KIC, for all $K \in \{U, C, W\}$, implies that KIC holds w.r.t. each probability distribution in $\mathcal{R}(P_{-i})$.[21] Our definition is easy to check.

As for i.i.d. types, we compare the set of social choice functions supportable under various notions of IC, when types evolve according to Markov processes.

Proposition 8 $\mathcal{F}^{MU}(\mu^*) \subset \mathcal{F}^{MC}(\mu^*) \subset \mathcal{F}^{MW}(\mu^*)$.

Since the one-shot game does not have a Markov counterpart, we cannot directly compare the sets of SCFs sustainable under dynamic and one-shot games.

[21] P_{-i} includes fewer probability distributions than $\triangle\Theta_{-i}$. Markov IC is thus weaker than incentive compatibility conditions for all possible probability distributions such as ex-post IC or Bayesian IC for all possible probability distributions (See Bergemann and Morris (2005) for details).

We now consider the case of private values, where agent i's type affects only the function u_i; with slight abuse of notation, the action profile α gives payoffs $u_i(\alpha, \theta_i)$ and $u_j(\alpha)$ respectively to agent i of type θ_i and principal j. While private values constitute a special case of interdependent values, we discuss it separately because it permits two simplifications. First, principal j' minmax values relative to complex mechanisms is the same for all notions of IC. Second, EDMs are not needed even off the equilibrium path; in fact constant mechanisms, i.e. actions, are enough to minmax any principal j. The simple minmax of principal j is

$$w_j := u_j(\alpha^j_{-j}, \alpha^j_j) = \max_{\alpha_j} u_j(\alpha^j_{-j}, \alpha_j), \text{ where } \alpha^j_{-j} \in \arg\min_{\alpha_{-j}\in\mathcal{A}_{-j}} \max_{\alpha_j\in\mathcal{A}_j} u_j(\alpha_{-j}, \alpha_j). \tag{15}$$

Proposition 9 shows that the principal's minmax value w^K_j is the same as the simple minmax w_j for any $K \in \{U, C, W\}$. Its proof will make use of the following.

Theorem 3 (Sion's Minmax Theorem) *Let X be a compact convex subset of a linear topological space and Y a convex subset of a linear topological space. If g is a real-valued function on $X \times Y$ with $g(x, \cdot)$ is upper semicontinuous and quasi-concave on Y for any $x \in X$, and $g(\cdot, y)$ is lower semicontinuous and quasi-convex on X for any $y \in Y$, then g has a saddle point $(\alpha, \beta) \in X \times Y$ where*

$$\max_{x\in X} \min_{y\in Y} g(x, y) = \min_{y\in Y} \max_{x\in X} g(x, y).$$

When principal j is being minmaxed, he is playing a best response but α^j_{-j} might not be the other principals' best responses, and it is necessary to ensure that they do not deviate. Since no principal directly observes mechanisms or actions chosen by other principals, he must rely on agents' reports. When j is being punished, principal $k \neq j$ offers a constant DDM ζ^j_k defined by

$$\zeta^j_k(\theta, d_k) = \alpha^j_k, \; \forall\theta, \; d_k = (d_{J+1,k} \ldots, d_{J+I,k}) \in \{0, 1 \ldots, J + I\}^I. \tag{16}$$

Proposition 9 *All minmax values are equal: $w^L_j = w_j \; \forall L \in \{1, U, C, W\}$.*

Proof The antecedents of Sion's Minmax Theorem 3 are satisfied with $X = \mathcal{A}_j$, $Y = \mathcal{A}_{-j}$, and $g = u_j$; therefore we have the saddle point $\alpha^j = (\alpha^j_{-j}, \alpha^j_j) \in \mathcal{A}$ where the maxmin and minmax are attained:

$$\min_{\alpha_{-j}\in\mathcal{A}_{-j}} \max_{\alpha_j\in\mathcal{A}_j} u_j(\alpha_{-j}, \alpha_j) = \max_{\alpha_j\in\mathcal{A}_j} \min_{\alpha_{-j}\in\mathcal{A}_{-j}} u_j(\alpha_{-j}, \alpha_j). \tag{17}$$

Step 1. We have $w_j^L \leq w_j$ because each principal $k \neq j$ can offer the constant DDM mechanism ζ_k^j that maps to only α_k^j, the mixed action profile that minmaxes j; given this, j can do no better than offer the constant DDM ζ_j^j mapping to α_j^j. Using the definition of u_j^L in (7),

$$\begin{aligned}
w_j^L := \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^L(\gamma_{-j},\gamma_j) &\leq \max_{\gamma_j\in\Gamma_j} u_j^L\left(\alpha_{-j}^j,\gamma_j\right) \\
&= \max_{\Gamma_j} \min_{\pi\in\Pi^L(\alpha_{-j}^j,\gamma_j)} \mathbb{E}_\mu\left[u_j\left(\pi\left(\theta\right)\right)\right] \\
&= \max_{\Gamma_j} \min_{\alpha_j\in\mathcal{A}_j(\gamma_j)} u_j\left(\alpha_{-j}^j,\alpha_j\right), \text{ since } u_j \text{ is indep. of } \theta \\
&\leq \max_{\Gamma_j} \max_{\alpha_j\in\mathcal{A}_j} u_j\left(\alpha_{-j}^j,\alpha_j\right) \\
&= \max_{\alpha_j\in\mathcal{A}_j} u_j\left(\alpha_{-j}^j,\alpha_j\right) = u_j(\alpha^j) = w_j.
\end{aligned}$$

Step 2. Now we show that $w_j^L \geq w_j$. Suppose j offers the constant DDM mapping to α_j^j. Given this, Sion's minmax theorem implies that the worst that principals other than j can choose for j is ζ_{-j}^j, or equivalently α_{-j}^j, even if principals other than j are able to use complex mechanisms to punish j:

$$\begin{aligned}
w_j^L := \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^L(\gamma_{-j},\gamma_j) &\geq \min_{\gamma_{-j}\in\Gamma_{-j}} u_j^L\left(\gamma_{-j},\alpha_j^j\right) \\
&= \min_{\alpha_{-j}\in\mathcal{A}_{-j}} u_j\left(\alpha_{-j},\alpha_j^j\right) = \max_{\alpha_j\in\mathcal{A}_j} \min_{\alpha_{-j}\in\mathcal{A}_{-j}} u_j(\alpha_{-j},\alpha_j) = w_j,
\end{aligned}$$

where the last equality is simply (17).

Steps 1 and 2 together imply $w_j^L = w_j$. Thus, principals offer $(\alpha_j^j, \zeta_{-j}^j)$ and the action profile played is $\alpha^j = (\alpha_{-j}^j, \alpha_j^j)$. ■

Two remarks are in order.

Remark Von Neumann's minmax theorem states that if X and Y are finite dimensional simplices and g is a bilinear function on $X \times Y$, then g has a saddle point (α, β):

$$\max_x \min_y \; g(x,y) = \min_y \max_x \; g(x,y) = g(\alpha,\beta).$$

If we were to apply this theorem then Y would be the set of all mixed strategies on A_{-j}, which is larger than $\mathcal{A}_{-j}$ since the former allows correlated punishments and the latter allows only independent mixing. Since repeated games usually restrict attention to independent, not just correlated, punishments we must use Sion's version rather than von Neumann's.[22]

[22] In the case of interdependent values, we cannot apply Sion's minmax theorem for w_j^{K*} defined in (11)

Remark The principal's minmax value w_j^1 in Yamashita's one-shot setting was not identified in terms of primitives in the case of interdependent values. Proposition 9 takes care of this. Under private values, principals' payoffs are independent of agents' types; this makes IC irrelevant when principal j is punished in both the one-shot game and the dynamic game and hence the principal's minmax value relative to complex mechanisms always equals the simple minmax w_j defined in (15) in both games: $w_j^U = w_j^C = w_j^W = w_j^1 = w_j$.

5.1 Folk Theorem under Private Values

5.1.1 I.I.D. Types and Private Values

Now we are ready to establish the folk theorem for the case of private values. What stage-SCFs can we support in PBE of $G^{\infty}(\delta)$ relative to Γ? We first consider the i.i.d. types with the product of independent distributions $\mu \in \triangle\Theta$. First we establish i.i.d folk theorems using CIC. As in the interdependent value case, the first folk theorem shows that any SCF in $\mathcal{F}^C(\mu)$ is supportable in a perfect Bayesian equilibrium of $G^{\infty}(\delta)$ relative to Γ, provided players are sufficiently patient. Similar to the case of interdependent values, principals can support any $f \in \mathcal{F}^C(\mu)$ by offering a profile of DDMs $\pi^a = (\pi_1^a, \ldots, \pi_J^a)$ such that

$$\pi_k^a(\theta, d_k) := f_k(\theta) \; \forall \, d_k \in \{0, 1, \ldots, J + I\}^J, \theta \in \Theta. \tag{18}$$

Agents then play a truthful continuation equilibrium in which agents report their true types and the identity of the deviating player (of course, 0 if no one deviates). A principal always believes that a majority of agents has reported truthfully. Note that a single agent's report does not impact the decision in a truthtelling equilibrium. In contrast to the case of interdependent values, the mechanisms offered by the non-deviating principals off the path following principal j's deviation are much simpler than EDMs due to Proposition 9. It shows that there is no loss of generality in assuming that all principals offer the profile of constant DDMs, $\zeta^j = (\zeta_1^j, \ldots, \zeta_J^j)$, where each ζ_k^j is specified by (16).

Proposition 10 *Under FD and private values, for any $f \in \mathcal{F}^C$ there exists $\underline{\delta} \in (0, 1)$ such that, for any $\delta \geq \underline{\delta}$, f is supportable as the equilibrium allocation of a perfect Bayesian equilibrium (PBE) of $G^{\infty}(\delta)$ relative to any complex Γ using only DDMs.*

Proof See the appendix. ■

Remark Similar to Proposition 4, we can weaken CIC - we can virtually support an SCF that satisfies WIC but not CIC.

We can establish the following propositions, similar to Propositions 5 and 6:

because the agent's incentive to report her type to the non-deviating principals depends on principal j's action.

Proposition 11 $\mathcal{F}_{+}^{1}(\mu) \subset \mathcal{F}^{U}(\mu) \subset \mathcal{F}^{C}(\mu) \subset \mathcal{F}^{W}(\mu)$.

Proposition 12 *If $f \in \mathcal{F}^{1}(\mu)$ then f can be supported in a PBE of the repeated game with low discounting using only DDMs.*

Unlike in the case of interdependent values, the principal's minmax values are the same across different notions of IC in the case of private values. The dynamic game supports more social choice functions in equilibrium solely because it weakens the notion of IC.

5.1.2 Markov Types and Private Values

If types follow independent Markov chains with the steady state distribution μ^*, we need to define the minmax values accordingly. As in the case of interdependent values, let MU, MC, MW denote Markov UIC, Markov CIC and WIC respectively. For each $K \in \{MU, MC, MW\}$, principal j's minmax value w_j^K is still equal to w_j according to Proposition 9. Consider agent i's minmax values without reservation payoffs. For $K \in \{MU, MC, MW\}$, w_i^K is equal to w_i specified in (10). With reservation payoffs, w_i^K is the same as agent i's reservation payoff specified in (8) with the expectation taken with the steady-state distribution μ^*.

Finally, let the stage-SCF f be SIR w.r.t μ^* with the proper minmax values given a notion of IC. For each $K \in \{MU, MC, MW\}$ define

$$\mathcal{F}^{MK}(\mu^*) := \{f \in \mathcal{F} \mid f \text{ is SIR w.r.t. } \mu^* \text{ and induced by } \pi \in \Pi^{MK}\}.$$

Proposition 13 *Let types follow an irreducible and aperiodic Markov chain in a model of private values where $u(\mathcal{F}^{MC}(\mu^*))$ satisfies FD. Any SCF $f \in \mathcal{F}^{MC}(\mu^*)$ is supported in a PBE of $G^{\infty}(\delta)$ for high δ relative to Γ, using only DDMs.*

Proof The proof follows from the Proposition 7 for interdependent values. ■

Remark Similar to Proposition 5.1.1, we can weaken CIC to virtually support an SCF that satisfies Markov WIC but not Markov CIC.

We now compare the set of social choice functions supportable under various notions of IC, when types evolve according to Markov processes. This proof is analogous to earlier propositions and hence omitted.

Proposition 14 $\mathcal{F}^{MU}(\mu^*) \subset \mathcal{F}^{MC}(\mu^*) \subset \mathcal{F}^{MW}(\mu^*)$.

Similar to Markov types for interdependent values, there is no corresponding one-shot game. However, Proposition 14 shows that the set of SCFs supported by equilibria in the dynamic setting expands as we adopt a weaker notions of Markov IC in the order of MU, MC, and MW.

6 Discussion and Conclusion

Our work on repeated decentralized contracting lies at the intersection of the literature on relational contracts, where continuation payoffs may be used to provide incentives for repeated short-term contracting, and the literature on competing mechanisms, where principals compete in the space of all contracts. The former literature provides simple mechanisms, but relies on there being only one principal. The latter literature incorporates competition among principals, but in a single-period setting; equilibrium payoffs are identified with reference to complex mechanisms.

Even static games with multiple principals offering competing mechanisms are beset with analytical difficulties. For intuition, consider a situation where multiple sellers compete in designing trading mechanisms. A seller cannot observe what trading mechanisms other sellers offer; or, even if he does, he cannot make his trading mechanism directly contingent on the others' mechanisms due to either lack of commitment power or institutional restrictions. Each buyer knows the market information, i.e. sellers' terms of trade (or trading mechanisms). If a seller can ask buyers to report both their types and their market information, he can make his terms of trade responsive to deviations by the competing sellers. This sustains equilibrium allocations that are unattainable with conventional direct mechanisms. However, it is a formidable task to design tractable mechanisms that are both useful for applications and can also support all possible equilibrium payoffs that one could with arbitrarily complex mechanisms.

Our main contribution is to offer an analysis of repeating contracting by multiple principals, opening the door to potential applications ranging from providing public goods to selling private goods, and from lobbying to financial contracting. When players are patient, we propose sufficient simple mechanisms. Principals offer DDMs (deviator-reporting direct mechanims) on the equilibrium path, and slight extention — EDMs — off path. These are much simpler than existing mechanisms, and fairly close to the DMs that are familiar from single-principal worlds. Our equivalence theorem shows that it is even simpler to characterize both the minmax and the equilibrium payoff set because the complex minmax of any principal in the repeated game can be expressed as his *maxmin* value when he can offer only actions, and the other principals are restricted to *DMs*.

Finally the dynamic setting supports more equilibrium payoffs than the one-shot setting does for two reasons. First of all, agents' endogenous monitoring completely neutralizes a deviating principal's ability to make his action choice contingent on agents' types off the path following his deviation. This is why a principal's minmax can be calculated as if he can only offer actions instead of complex mechanisms. Secondly, the dynamic setting allows weaker notions of IC. To show this, we establish various notions of IC — UIC, CIC, and WIC, from the strongest to the weakest. UIC deters all possible false messages, and is the right notion for the static setting. The other two notions check incentive compatibility against a smaller class of feasible deviations. We propose a notion of consistency of deviations that is useful in separating deviations that must be deterred with contemporaneous incentives, and those that can be punished in the continuation game. This allows us to use CIC or WIC in the dynamic setting instead of UIC.

Some assumptions of our model merit discussion. First, we use the fact that there are three agents to check reports, so that unilateral deviations from truthful reporting on the identity of a deviating player are always detected. However note that our observability assumption is very weak, in that principals have no information except that from agents. If we allow principals to observe (with a time lag) the mechanisms offered by the other principals and the actions taken by them, our results immediately extend to any number of agents. We conjecture that it is possible to extend our results to two agents even if principals have imperfect statistical information.

Principals need not ask agents to report the identity of a deviator every period. When players are sufficiently patient, principals only need to ask, for example, every ten periods or every twenty periods who the last agent to deviate was. What is the minimum number of principals that need to know the deviation to punish it? This depends on the environment. If principals are sellers and agents are buyers, it may be possible that only one principal need to know the deviation especially if it is an agent's deviation.

A Appendix

Proof of Lemma 2 Fix $\gamma_{-j} \in \Gamma_{-j}$. For any given α_j, let $\Pi^K_{-j}(\gamma_{-j}, \alpha_j) := \Pi^K_{-j}(\alpha_j) \cap \Pi_{-j}(\gamma_{-j}, \alpha_j)$. Using the definition of u^K_j from (7), for any $\gamma_j \in \Gamma_j$ we have

$$\begin{aligned}
\max_{\alpha_j \in \mathcal{A}_j} u^K_j(\gamma_{-j}, \alpha_j) &= \max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi^K_{-j}(\gamma_{-j}, \alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta) \right] \\
&\geq \min_{\alpha_j \in \mathcal{A}_j(\gamma_j)} \min_{\pi_{-j} \in \Pi^K_{-j}(\gamma_{-j}, \alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta) \right] \\
&\geq \min_{\pi \in \Pi^K(\gamma)} \mathbb{E}_\mu \left[u_j(\pi(\theta), \theta) \right] \\
&=: u^K_j(\gamma_{-j}, \gamma_j).
\end{aligned}$$

Taking max over all Γ_j, we have

$$\max_{\alpha_j \in \mathcal{A}_j} u^K_j(\gamma_{-j}, \alpha_j) \geq \max_{\gamma_j \in \Gamma_j} u^K_j(\gamma_{-j}, \gamma_j).$$

The reverse inequality follows immediately since Γ_j includes the set of all constant mechanisms. This proves the lemma. ■

Proof Proposition 1 Step 1. We have $w_j^K \leq w_j^{K*}$ because principals $k \neq j$ can offer λ_k^{Kj}; given this, Lemma 2 implies that principal j can do no better than play the constant mechanism that always assigns α_j^{Kj}, even if principal j is able to use complex mechanisms.

$$w_j^K = \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^K(\gamma_{-j},\gamma_j) = \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\alpha_j\in\mathcal{A}_j} u_j^K(\gamma_{-j},\alpha_j) \leq \max_{\alpha_j\in\mathcal{A}_j} u_j(\lambda_{-j}^{Kj},\alpha_j) = w_j^{K*},$$

where $u_j(\lambda_{-j}^{Kj},\alpha_j)$ is principal j's expected payoff based on the truthful continuation equilibrium.[23] The second equality follows from Lemma 2. The inequality holds because λ_{-j}^{Kj} is one profile of mechanisms available for principals except for j.

Step 2. Now we show that $w_j^K \geq w_j^{K*}$. The first inequality below holds because the constant mechanism that always assigns α_j^{Kj} is one of mechanisms available for principal j. The second equality holds because any continuation equilibrium at $(\gamma_{-j},\alpha_j^{Kj})$ generates DMs π_{-j} for principals except for j that satisfy KIC conditional on α_j^{Kj}.

$$\begin{aligned} w_j^K &:= \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^K(\gamma_{-j},\gamma_j) \geq \min_{\gamma_{-j}\in\Gamma_{-j}} u_j^K(\gamma_{-j},\alpha_j^j) \\ &= \min_{\pi_{-j}\in\Pi_{-j}^K(\alpha_j^j)} \mathbb{E}_\mu\left[u_j(\pi_{-j}(\theta),\alpha_j^j,\theta)\right] = \max_{\alpha_j\in\mathcal{A}_j} u_j(\lambda_{-j}^{Kj},\alpha_j) = w_j^*. \end{aligned}$$

Steps 1 and 2 together imply $w_j^* = w_j^{K*}$. ∎

Proof Proposition 2 Step 1: We show that $w_j^U \leq w_j^1$. The first equality below holds by Lemma 2. The second equality follows from the definition of $u_j^U(\gamma_{-j},\alpha_j)$. The first inequality holds because $\Pi_{-j}^U(\gamma_{-j},\alpha_j) \supset \Pi_{-j}^1(\gamma_{-j},\alpha_j)$, where $\Pi_{-j}^1(\gamma_{-j},\alpha_j)$ is the set of all profiles of the other principals' UIC DMs that are induced by all continuation equilibria at (γ_{-j},α_j) in the static game. The third equality follows from the definition of $u_j^1(\gamma_{-j},\alpha_j)$. The last inequality holds because $\mathcal{A}_j$ is equivalent to the set of all constant mechanisms in Γ_j and is a subset of Γ_j.

$$\begin{aligned} w_j^U &:= \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\alpha_j\in\mathcal{A}_j} u_j^U(\gamma_{-j},\alpha_j) \\ &= \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\alpha_j\in\mathcal{A}_j} \min_{\pi_{-j}\in\Pi_{-j}^U(\gamma_{-j},\alpha_j)} \mathbb{E}_\mu\left[u_j\left(\pi_{-j}\left(\theta\right),\alpha_j,\theta\right)\right] \\ &\leq \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\alpha_j\in\mathcal{A}_j} \min_{\pi_{-j}\in\Pi_{-j}^1(\gamma_{-j},\alpha_j)} \mathbb{E}_\mu\left[u_j\left(\pi_{-j}\left(\theta\right),\alpha_j,\theta\right)\right] \\ &= \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\alpha_j\in\mathcal{A}_j} u_j^1(\gamma_{-j},\alpha_j) \\ &\leq \min_{\gamma_{-j}\in\Gamma_{-j}} \max_{\gamma_j\in\Gamma_j} u_j^1(\gamma_{-j},\gamma_j) \ =: w_i^1 \end{aligned}$$

Step 2. Because $w_j^K = w_j^{K*}$ for all $K \in \{U, C, W\}$ by Proposition 1, the first inequality above

[23] Proposition 3 shows how truthful reporting is enforced in the repeated game.

follows if we can show that $w_j^{W*} \leq w_j^{C*}$ and $w_j^{C*} \leq w_j^{U*}$. Inequalities (19) and (20) hold because $\Pi_{-j}^W(\alpha_j) \supset \Pi_{-j}^C(\alpha_j)$ and $\Pi_{-j}^C(\alpha_j) \supset \Pi_{-j}^U(\alpha_j)$ for all α_j:

$$\max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi_{-j}^W(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta)\right] \leq \max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi_{-j}^C(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta)\right]; \tag{19}$$

$$\max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi_{-j}^C(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta)\right] \leq \max_{\alpha_j \in \mathcal{A}_j} \min_{\pi_{-j} \in \Pi_{-j}^U(\alpha_j)} \mathbb{E}_\mu \left[u_j(\pi_{-j}(\theta), \alpha_j, \theta)\right]. \tag{20}$$

From steps 1 and 2, we have $w_j^W \leq w_j^C \leq w_j^U \leq w_j^1$. ■

Proof of Lemma 3 Given full-dimensionality we can construct[24] a PSP $\{\overline{\beta}^i \mid i \in \mathcal{N}\}$. Since $\overline{\beta}^i \in u(\mathcal{F}^C(\mu))$, by construction there exists a family of DDMs $\{\overline{\pi}^i \mid i \in \mathcal{N}\}$ such that $\overline{\beta}^i := \mathbb{E}_\mu u\left(\overline{\pi}^i(\theta)\right)$. Since properties 1,2 and 3 above rely on strict inequalities it is easy to see that there is an $r > 0$ such that any family $\{y^i \mid y^i \in B(\overline{\beta}^i, r)\}$ is also a PSP. If any such $\overline{\pi}_k^i$, for $i \in \mathcal{N}$ and $k \in \mathcal{J}$, is not one-to-one, replace it with a new DDM π_k^{ai} as follows. (If any $\overline{\pi}_k^i$ is one-to-one, set $\pi_k^{ai}(\theta) = \overline{\pi}_k^i(\theta)$.) Fix any enumeration of the type-space $\Theta = \{\theta^1, \ldots \theta^L\}$. Let

$$\pi_k^{ai}(\theta^1) := \overline{\pi}_k^i(\theta^1),$$

and for $l \geq 1$ pick an arbitrary element

$$\pi_k^{ai}(\theta^{l+1}) \in B\left(\overline{\pi}_k^i(\theta^l), r\right) \setminus \left\{\pi_k^{ai}(\theta^1), \ldots, \pi_k^{ai}(\theta^l)\right\}.$$

Now define $\beta^i := \mathbb{E}_\mu[u\left(\pi^{ai}(\theta)\right)]$ for all $i \in \mathcal{N}$. ■

Proof of Proposition 3 Fix any $f \in \mathcal{F}^C(\mu)$ and let $v := u(f)$ denote the target payoff. Strategies are defined by the following rules.

1. Each principal k starts off the game in Phase I, when he offers the DDM π_k^a comprising (i) a message space $E_{ik} := \{0, 1 \ldots, J, J+1, \ldots, J+I\} \times \tilde{\Theta}_i$ for each agent i, and (ii) a mapping
$$\pi_k^a : E_k \to \mathcal{A}_k, \text{ where } E_k = \times_{i \in I} E_{ik}.$$
Each π_k^a induces the corresponding f_k if agents report truthfully:
$$\pi_k^a(\theta, (d_{ik})_i) = f_k(\theta) \text{ for all } d_{ik} \in \{0, 1 \ldots, J, J+1, \ldots, J+I\}^I.$$

2. Agents start in phase I, reporting $(\theta_i^t, 0)$ to all principals at time t. If principal j deviates unilaterally (offers a contract other than π_j^a), agents play myopic best responses in the current period and send messages $e_{ik} = (m_{ik}(\theta_i^t), j)$; here $m_{ik}(\theta_i)$ is the message sent

[24] See Abreu, Dutta and Smith (1991).

from i to k in any *equilibrium* of the static game. Similarly, if agent i deviates. If it is clear that an agent deviated at time t but the identity of the deviating agent is not clear, report $d_{ik}^{t+1} = J + 1$. If a majority of agents reports j then switch to Phase II^j; if the majority report in each period of Phase II^j is $d_{ik}^t = 0$, then switch to Phase III^j with probability q.

3. When principal $k \neq j$ receives the report $d_{ik}^t = j \in \mathcal{J}$ from a majority of agents, he offers the EDM defined by equation (13) in phase II^j; however, principal j can potentially deviate to a complex mechanism; we specify a mapping $g^j : \Gamma_j \to \mathcal{A}_j$ defined by (12) that describes the action of j agents will induce from each γ_j regardless of their types. Suppose that principal j offers a complex mechanism γ_j^s in period s during her punishment phase. Then, agents are asked to induce the action $g^j(\gamma_j^s)$ regardless of their types. If the profile of messages $m_j^* = (m_{ij}^*)_{i \in \mathcal{I}}$ induces $g^j(\gamma_j^s)$, that is, $\gamma_j(m_j^*) = g^j(\gamma_j^s)$, then each agent i sends m_{ij}^* to principal j regardless of her type. At the time of reporting to principal k, agents do not observe principal j's action but each agent i is asked to report $p_{ikj}^t = g^j(\gamma_j^t)$ to each principal k $(k \neq j)$ expecting that $g^j(\gamma_j^t)$ will be induced. If principal j' deviates from a punishment phase at t, each agent i reports $d_{ik}^t = j'$.

4. When principal k receives the report $d_{ik}^t = i \in \mathcal{I}$ from a majority of agents, he moves to phase II^i, where he offers the DMs π^i that attain the minmax value w_i^K defined in equation (9); Lemma 1 ensures that the minmax can be reached with a profile of DMs.

5. In Phase III^j play $\pi^{ai} \in \Pi$, giving the payoff vector β^j.

6. One of the key differences between this and the usual repeated games proofs is that we need to consider histories where a principal j who has actually not deviated has been reported as the deviator by a majority of agents to the other principals, or a deviating principal has not been reported to j by a majority. At such histories j unexpectedly finds himself being minmaxed or not minmaxing a player who had deviated in the previous period. In such a case j proceeds as if he had indeed deviated in the current period; agents, including those who misreported earlier, behave as if j had indeed deviated and subsequently report him as the deviator. Such histories do not happen as a result of unilateral deviations or on the equilibrium path.

Choice of Parameter: The only parameter in the above strategy is the probability q. Let $\max_{\mathcal{I},A,\Theta} |u_i(a,\theta)| < \bar{M} < \infty$. Pick any $q \in (0,1)$ such that

$$\bar{M}(1-q) < \beta_i^i(2-q) - w_i^C \; \forall i \in \mathcal{I} \cup \mathcal{J}. \tag{21}$$

Note that at $q = 1$ this inequality becomes $0 < \beta_i^i - w_i^C$.

Verification of Equilibrium: We show that the proposed strategy is unimprovable, i.e. no one-shot deviation by any player from any phase is profitable. Since f satisfies CIC, all profitable deviations from it are observable with positive probability; let

$$p_{\min} := \min_{(\tilde{\theta}_{ij})_j \in (\tilde{\Theta})^J} \mu \left\{ \theta_{-i} \in \Theta_{-i} \mid (\pi_1(\tilde{\theta}_{i1}, \theta_{-i}), \ldots, \pi_J(\tilde{\theta}_{iJ}, \theta_{-i})) \notin \hat{\mathcal{A}}(\pi) \right\}$$

be the minimum probability of detection after a (strictly) profitable deviation. All deviations of consequence in phases II and III may be detected w.p. 1. The quantity $p_{\min}$ plays a role only in Phase I.

1. Phase II^i: Player i's "lifetime" (discounted average) payoff in phase II^i, denoted L_i^i, satisfies $L_i^i = (1-\delta)w_i^C + \delta\left(qL_i^i + (1-q)\beta_i^i\right)$ so that

$$L_i^i = \frac{(1-\delta)w_i^C + \delta(1-q)\beta_i^i}{1-\delta q}. \tag{22}$$

   Note that $L_i^i \to \beta_i^i$ as $\delta \to 1$. Player i will not deviate since the maximal payoff to a one-shot deviation is below $(1-\delta)w_i^C + \delta L_i^i$, which is lower than L_i^i. Player $j \neq i$ will not deviate for high δ, since his maximal payoffs are bounded above by $(1-\delta)\bar{M} + \delta L_j^j$, which tends towards β_j^j, while his payoff L_j^i from conformity tends to $\beta_j^i > \beta_j^j$ (payoff asymmetry).

2. Phase III^i: Note that the mechanisms used in this phase are all one-to-one by construction; this means that any deviator, not just principals, is common knowledge among agents, and the quantity $p_{\min}$ plays no role. From the definitions it is clear that the difference in the lifetime payoffs to one-shot deviation and conformity is bounded above by

$$(1-\delta)\bar{M} + \delta L_i^i - \beta_i^i = (1-\delta)\left[\bar{M} - \frac{(1+\delta-\delta q)\,\beta_i^i - \delta w_i^C}{1-\delta q}\right] \tag{23}$$

   using (22). An immediate implication of inequality (21) defining q is that (23) is strictly negative for all δ close to 1. Since $\beta_j^i > \beta_j^j\, \forall j \neq i$, it is immediate that players $j \neq i$ do not have a profitable one-shot deviation either.

3. Phase I: If player i deviates from Phase I, he/she earns a continuation payoff of L_i^i in the repeated game with probability at least $p_{\min}$. Note that unlike a usual repeated game this probability may be strictly lower than 1 in case an agent sends an inconsistent message outside her B_i. (Recall that inconsistent messages within B_i cannot be detected.) Since $v_i > \beta_i^i$ (target payoff domination) the arguments above also imply that $(1-\delta)\bar{M} + \delta\{p_{\min}L_i^i + (1-p_{\min})v_i\} < v_i$, for high δ; the strategies in phase I are therefore unimprovable.

4. Agent $i \in \mathcal{I}$ reports (p^t_{ikj}, d^t_{ik}) truthfully: since $I \geq 3$ and only the majority's report matters, no agent can change the continuation game, either the current action of a principal or the phase transition, by unilaterally changing one or both of the first two components of his report. His type report can influence each principal's current action, but since $\psi^{Kj}_k(\alpha^s_j)$ is CIC over agents' types, it implies that λ^{Kj}_k induces agents to report truthfully within the set B_i; a report outside B_i leads to an unexpected action and is therefore deterred in the steps above.

In sum, for high δ, the posited strategy is unimprovable after all histories, and hence is an equilibrium. ■

Proof Proposition 4 Take $\pi \in \mathcal{F}^1(\mu)$. Fix $\varepsilon > 0$. In addition to being WIC, this satisfies SIR for all principals and WIR for all agents. From Abreu and Matsushima we know that for each $i \in \mathcal{I}$ there is a mapping $h^i : \Theta_i \to \triangle A$ such that

$$u_i(h^i(\theta_i), \theta_i) > u_i(h^i(\theta'_i), \theta_i) \quad \forall \theta'_i \neq \theta_i.$$

In other words, these lotteries on A give a strict incentive to report truthfully. Define a new DM that satisfies *strict* IC over all consistent messages:

$$\pi^+(\theta) := (1-\Delta)\pi(\theta) + \frac{\Delta}{I}\sum_{i\in\mathcal{I}} h^i(\theta_i).$$

Pick $\Delta > 0$ small enough that $\parallel \pi - \pi^+ \parallel < \varepsilon/2$. Now we perturb π^+ to an invertible DM π^* such that $\parallel \pi^+ - \pi^* \parallel < \varepsilon/2$. First, fix an enumeration of the type-space $\Theta = \{\theta^1, \ldots, \theta^L\}$. By definition there is $r \in (0, \varepsilon/2JL)$ such that $B(\pi^+_j(\theta^l), r) \cap \mathcal{F}^W(\mu) \neq \emptyset$ for any $j \in \mathcal{J}$. Let $\pi^*_j(\theta^1) := \pi^+_j(\theta^1)$; and pick an arbitrary element

$$\pi^*_j(\theta^{l+1}) \in B\left(\pi^+_j(\theta^l), r\right) \setminus \left\{\pi^*_j(\theta^1), \ldots, \pi^*_j(\theta^l)\right\} \text{ for } l \geq 1.$$

By construction we therefore have $\parallel \pi^* - \pi \parallel < \varepsilon$. Since each π^*_j is invertible, $B_i(\pi^*) = \{(\theta_i, \ldots, \theta_i) \in (\Theta_i)^J \mid \theta_i \in \Theta_i\}$ for any agent i; consequently, an inconsistent report by agent i is immediately detected and leads to her being punished. By Proposition 3 it follows that there is a $\underline{\delta} \in (0, 1)$ such that π^* can be supported in a PBE of the repeated game. ■

Proof of Proposition 7 Let the payoff vector from the SCF f be $v = (v_1, \cdots, v_{I+J})$. Using Lemma 3 construct the PSP for v. In what follows i and j are generic players unless otherwise identified.

Equilibrium Strategies:

Phase I is the same as before: Each principal offers a DDM, and throughout believes whatever the majority of agents report as the identity of the deviating principal. Phase II^j

is the same except that in the latter phase we minmax with probability 1 for $n+m$ periods (the equilibrium value of m and n will be chosen subsequently) rather than a single period, and then decide to continue minmaxing for the next block of $m+n$ period with probability $q \in (0,1)$. This minmaxing is done by playing the profile of EDMs λ_{-j}^{Kj}. If i unilaterally deviates, agents report it to the other principals. Subsequently, the game moves to phase II^i. Else, with probability q stay in phase II^j, while with probability $(1-q)$ proceed to phase III^j. When in phase III^j play the action profile that gives the payoff vector β^j in the *stationary distribution*. Stay here unless i deviates unilaterally and takes the game to phase II^i.

Suppose all principals play player i's minmax profile for $n+m$ periods. During the first n periods the probability distribution over the agent's types could be very different from μ_i^* and the agent's expected payoff could be as high as $\bar{M}$. Let $L = \#\Theta$. Using Theorem 2, pick n is large enough that for all $i \in \mathcal{I} \cup \mathcal{J}, \theta_i \in \Theta_i$,

$$\left|\mathbb{P}(\mathbb{X}_i^t = \theta) - \mu_i^*(\theta_i)\right| < \epsilon/3L\bar{M} \ \forall t \geq n.$$

This means that in each of the next m periods his expected mean (undiscounted) payoff is within $\epsilon/3$ of u_i^*. Now pick m large enough that

$$\left|\frac{n\bar{M} + m(w_i^K + \epsilon/3)}{n+m} - w_i^K\right| < \frac{2\epsilon}{3}$$

There exists a discount factor $\delta^\dagger(n,m) \in (0,1)$ such that $\bar{M}(1-\delta)\sum_{t=0}^{n+m-1}\delta^t < \epsilon/3$ for all $\delta \geq \delta^\dagger(n,m)$. Player i's δ-discounted payoff over $n+m$ periods is within $\epsilon/3$ of the undiscounted payoff. Let $\bar{u}_i$ denote the average δ-discounted payoff over the $n+m$ periods in phase II^j. Our choice of n, m, $\delta^\dagger(n,m)$ ensures that $\bar{u}_i$ is within ϵ of w_i^C if $\delta \geq \delta^\dagger(n,m)$.

Verification of equilibrium: This is similar to the earlier proposition and hence omitted. ■

Proof of Proposition 10 The target (expected) payoff vector is $v \equiv (v_1, \ldots, v_{J+I}) := \mathbb{E}_\mu u(f)$. Take β^is and π^{ai}s as in Lemma 3 for each $i \in \mathcal{I} \cup \mathcal{J}$.

Equilibrium Strategies. The strategy vector that generates the target payoff v as an equilibrium payoff (for suitably chosen parameters) is defined in Markov strategy terminology as follows.

1. Start in phase I, where principals offer the deviator-reporting DMs $\pi^a = (\pi_1^a, \ldots, \pi_J^a)$ satisfying (18).

   As long as no principal deviated in the current play of phase I and no agent deviated in the last period, each agent $i \in \mathcal{I}$ reports $(\theta_i^t, 0)$ to each principal at t, where θ_i^t is her true type; this induces the vector of actions $f(\theta)$ as a function of the agents' types.

Otherwise the reports are (θ_i^t, d_{ik}^t) with $d_{ik}^t \neq 0$. If π^t is the mechanism offered at time t, the action profile $\alpha^t \notin \hat{\mathcal{A}}(\pi^{at})$ only if (at least one) an agent had deviated. If it is clear that agent $i \in \mathcal{I}$ deviated at $t-1$, set all $d_{ik}^t = i$. If the identity of the deviating agent is not clear, $d_{ik}^t = J+1$. If $d_{ik}^t = i$ go to Phase II^i.

If $j \in \mathcal{J}$ deviates unilaterally (offers a contract other than π_j^a), agents play any equilibrium of the one-shot game induced by the given mechanisms and report (θ_i^t, j) to all the other principals. Each principal who offers a DDM believes that if more than half the agents agree that j is the deviating principal, he is indeed guilty; principal k then moves to phase II^j, and stays in phase I otherwise. For any profile of reports d_k received by principal k, let $\hat{d}_k$ denote the majority report [25] from $\{d_{ik} : i \in \mathcal{I}\}$. In particular, if one or more agents deviate, a principal who offers the right mechanism in Phase I may nevertheless play an unexpected action; in such a situation the pricipal is not considered as a deviator and agents are punished as described above.

2. When in phase II^j for some $j \in \mathcal{J}$, each $k \in \mathcal{J}$ offers the constant deviator-reporting DDM ζ_k^j. If any i unilaterally deviates, agents report it. Subsequently, the game moves to phase II^i. Else, with probability q stay in phase II^j, while with probability $(1-q)$ proceed to phase III^j.[26] In phase II^i for some $i \in \mathcal{I}$, principal k plays the action α_k^i that attain the minmax value $w_i^K = w_i$ as in equation (10).

3. When in phase III^j the action profile that gives the payoff vector β^j is played. Stay here unless $i \in \mathcal{I} \cup \mathcal{J}$ deviates unilaterally and takes the game to phase II^i.

4. One of the key differences between this and the usual repeated games proofs is that we need to consider histories where a principal j who has actually not deviated has been reported as the deviator by a majority of agents, or a deviating principal has not been reported. At such histories j unexpectedly finds himself being minmaxed or not minmaxing a player who had deviated in the previous period. In such a case j proceeds as if he had indeed deviated in the current period; agents, including those who misreported earlier, behave as if j had indeed deviated and report him again as the deviator. Such histories do not happen as a result of unilateral deviations or on the equilibrium path.

In words, the strategy says: Start and continue with the deviator-reporting DMs given by f until the first unilateral deviation by a principal (say by i). Then, minimax i for one period (with probability one) and (in the event of no observed deviation) continue the minimaxing with probability q. With the remaining probability, terminate the minimaxing and play III^i until further deviations. Treat players symmetrically and subject every unilateral deviation

[25] If the majority is not unique pick the smallest integer among all the candidates.

[26] This coordination is done through the use of a PCD as usual; so there is no possibility of miscoordination.

to this (stochastic) punishment schedule. Multilateral deviations are ignored, as in any Nash equilibrium. Along the lines of earlier proofs it may be checked that this is sufficient to deter deviations when players are patient. ■

References

[1] Abreu, D., D. Dutta and L. Smith (1991): "The Folk Theorem for Repeated Games: A NEU Condition," *Econometrica*, 62(4), 939-948.

[2] Abreu, D. and H. Matsushima (1992): "Virtual Implementation in Iteratively Undominated Strategies I: Complete Information," *Econometrica*, 60, 993-1008.

[3] Attar, A., Mariotti T., and F. Salanié (2011): "Non-Exclusive Competition in the Market for Lemons," *Econometrica*, forthcoming.

[4] Baker, G.P., R. Gibbons, and K.J. Murphy (2002): "Relational Contracts and the Theory of the Firm," Quarterly Journal of Economics, 117, 39-84.

[5] Bergemann, D. and J. Välimäki (2003): "Dynamic Common Agency," *Journal of Economic Theory*, 111(1), 23-48.

[6] Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73(6), 1771–1813.

[7] Board, S. (2011):"Relational Contracts and the Value of Loyalty," *American Economic Review,* 101(7), 3349–67.

[8] Breiman, L.: Probability, Society for Industrial and Applied Mathematics, 1991.

[9] Chandrasekher, M. (2012) Unraveling in a Repeated Moral Hazard Model with Multiple Agents, ASU mimeo.

[10] Crémer, J. and R. McLean(1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

[11] Epstein, L. and M. Peters (1999): "A Revelation Principle for Competing Mechanisms," *Journal of Economic Theory*, 88(1), 119-160.

[12] Evans, R. (2008) Simple Efficient Contracts in Complex Environments, Econometrica....

[13] Fong, Y.-F., and J. Li (2012): "Relational Contracts, Efficiency Wages, and Employment Dynamics," mimeo.

[14] Han, S. (2006): “Menu Theorems for Bilateral Contracting,” *Journal of Economic Theory*, 131 (1), 157-178.

[15] Lee, J. and H. Sabourian (2011): “Efficient Repeated Implementation,” *Econometrica*, 79(6), 1967-1994.

[16] Levin, J. (2003):”Relational Incentive Contracts,” *American Economic Review*, 93, 835-857.

[17] Ma, C. (1988): \Unique Implementation of Incentive Contracts with Many Agents," The Review of Economic Studies, 55, 555-571.

[18] W. Bentley MacLeod and James M. Malcomson (1993) “Investments, Holdup, and the Form of Market Contracts,” The American Economic Review Vol. 83, No. 4, pp. 811-837.

[19] Martimort, D. and L. Stole (2002): “The Revelation and Delegation Principles in Common Agency Games,” *Econometrica*, 70(4), 1659-1673.

[20] Matsushima, H. (1988): “A New Approach to the Implementation Problem,” *Journal of Economic Theory*, 45(1), 128-144.

[21] McAfee, P. (1993), “Mechanism Design by Competing Sellers,” *Econometrica*, 61, 1283-1312.

[22] Obara, I. and D. Rahman (2010), “Mediated Partnerships,” Econometrica 78, 285-308.

[23] Pavan, A. and G. Calzolari (2010): “Truthful Revelation Mechanisms for Simultaneous Common Agency Games,” *American Economic Journal: Microeconomics*, 2(2), 132-190.

[24] Pavan, A., I. Segal, and J. Toikka (2011): “Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization Problem, and Information Disclosure,” mimeo, Northwestern University.

[25] Peck, J. (1995): “Competing Mechanisms and the Revelation Principle,” Ohio State University, unpublished manuscript.

[26] Peters, M. (2001): “Common Agency and the Revelation Principle,” *Econometrica*, 69(5), 1349-1372.

[27] Peters, M. and B. Szentes (2011), “Definable and Contractable Contracts,” *Econometrica*, forthcoming

[28] Prat, A. and A. Rustichini (2003), "Games Played through Agents," *Econometrica,* 71(4), 989–1026.

[29] Salanie and Rey....

[30] Sion, M. (1958): "On General Minimax Theorems," *Pacific Journal of Mathematics*, 8(1), 171-176.

[31] Tripathi, V. (2008). "Essays on Strategic Voting and Mechanism Design," Princeton University thesis.

[32] Yamashita, T. (2010), "Mechanism Games with Multiple Principals and Three or More Agents," *Econometrica*, 78(2), 791-801.