Contractible Contracts in Common Agency Problems

Balázs Szentes*

December 16, 2013

Abstract

This paper analyzes contractual situations between many principals and many agents. The agents have private information, and the principals take actions. Principals have the ability to contract not only on the reports of the agents but also on the contracts offered by other principals. Contracts are required to be representable in a formal language. The main result of the paper is a characterization of the allocations that can be implemented as equilibria in our contracting game. When we restrict attention to exclusive-contracting environment, our characterization result implies that principals can collude to implement the monopolist outcome. Finally, in general, equilibrium contracts turn out to be *incomplete*. That is, a contract will restrict the action space of a principal but will not necessarily determine a single action.

1 Introduction

In many settings, firms do not charge a fixed price, but instead make their prices explicitly conditional on the prices offered by competitors in an apparent effort to attract customers. They often commit to *price relationship agreements*, that is, they adopt policies which are directly linked to the price policies of other firms. Examples of these policies include meet-the-competition clauses, price-beating promises and lowest fare guarantees. It is not immediately clear whether these policies are indeed beneficial to consumers, or whether they simply enable collusion between firms. We are therefore motivated to explore in more general terms the possibility of contracts which depend on the contracts offered by others.[1] Indeed, our goal in this paper is to put forward a general common-agency model and then explore the consequences of allowing contractibility of contracts.

*Department of Economics, London School of Economics.

[1] Reciprocal trade agreements, such as GATT, also take the form of contractible contracts. A reciprocal contract commits to setting a low tariff against a particular country if that country's contract does the same. Finally, tax treaties sometimes have this flavor – for example, out of state residents who work in Pennsylvania are exempt from paying Pennsylvania tax as long as they live in a state that has a reciprocal agreement exempting Pennsylvania residents from state taxes. See http://www.revenue.state.pa.us/revenue/cwp/view.asp?A=238&Q=244681.

The contribution of this paper is threefold. First, we show that by allowing for contractible-contracts we are able to provide a full characterization of implementable allocations. Indeed, we prove a *folk theorem*. The ability to contract on contracts gives principals the opportunity to collude, and thereby implement a variety of outcomes, as in repeated games. Collusion is accomplished through contracts which punish a principal if his contract is not the one expected from him in equilibrium. Since contracts are contractible, a principal is able to commit to punishing a deviator despite the fact that interactions are not repeated. Our characterization theorem is presented in greater detail below.

Second, we investigate the implications of our general results when applied to *exclusive-contracting* environments (where an agent can sign a contract with only one principal). This type of exclusivity exists, for example, in the context of employer-worker relationships in which the worker may accept only one job offer, and seller-buyer relationships in which the buyer is interested in purchasing a good or a service from only one seller. If contracts were not contractible, the competition among principals would result in a Pareto efficient outcome. In contrast, we show that the contractibility of contracts enables principals to collude and effectively act as a monopolist, offsetting any efficiency gain generated by competition. Therefore, in such environments, the prohibition of contracting on contracts emerges as a policy implication.

Third, our theory provides a rationale for the ubiquity of *incomplete contracts*. A contract is referred to as incomplete if some terms of the agreement between the principal and the agent are at the principal's discretion even after the contract is signed. In our model, whether or not we allow for the contractibility of contracts, equilibrium contracts will, in general, be incomplete. Incomplete contracts are common in many contexts. For example, labor contracts often specify a fixed wage but allow the bonus to be at the employer's discretion, adjustable rate mortgages permit the lender to change the interest rate within pre-specified bounds, and utility and other long-term service providers may unilaterally raise prices.

In the specific model analyzed in this paper, there are several principals and several agents. Agents have types, and principals take actions. Each principal wishes to enter into a contractual relationship with each agent. Following the usual approach in the literature, we analyze equilibria in *communication games*. In a communication game, agents are endowed with message spaces, and the game has three stages. At the first stage, the principals offer contracts to the agents simultaneously. In our setting, a contract is a mapping from the messages of an agent and contract profiles of the principals to the subsets of the action space of the principal.[2] During the second stage, agents simultaneously send private messages to each principal. In the final stage, principals select actions from the subsets of actions determined by both the first-stage contracts and the second-stage message profiles. Our goal is to characterize the set of equilibrium outcomes of these

[2] In existing literature on common agency models it is usually assumed that contracts are complete, that is, they determine a single action for the principal as opposed to a subset. We show that this assumption results in a loss of generality.

games.

Such models give rise to an *infinite regress problem*, which can make them difficult to solve. Consider two principals whose payoffs both depend on the action of the other. Each principal's action is determined by his contract with the agents, so each principal would like to offer a contract which is contingent on the contract offered by the other principal. A principal's contract will typically be contingent on the contract of the other principal, which, in turn, is contingent on the contract of the first principal, and so on. One possible way of dealing with this hierarchical dependency is to include *self-referential* contracts in the contract space. In a self-referential contract, the action a principal takes will depend on whether the other principal offers the same self-referential contract. That is, the contract refers to itself. The use of this type of contract allows us to collapse many statements about higher order dependencies into a single self-referential statement. However, the construction of a contract space which includes such contracts is not immediately obvious. Therefore, perhaps the most important feature of a common agency model is the set of contracts available to the principals. Following the approach of Peters and Szentes (2012), this paper models the space of contracts with the set of *definable correspondences*. This set is a generalization of recursive functions, and will be discussed in detail in the next section. As will be shown, the key feature of this space is that it includes all sorts of self-referential mappings.

Our main result consists of a folk theorem, which asserts that an allocation can be implemented as an equilibrium in our contracting game if and only if the allocation is *subgame-implementable* and the induced payoff of each principal is larger than his minmax value. We shall provide an explanation for both subgame implementability and the minmax value below. In order to do so, we first define an *ordinary* contract to be one that does not condition on the contracts of the other principals.[3] Consider a modification of our contracting game in which each principal must offer an ordinary contract at the first stage. An allocation is called subgame-implementable if it is an equilibrium outcome in a subgame generated by *some* ordinary contract profile. Let us now provide a clear definition of Principal q's minmax value. Suppose that for each $j \neq q$, Principal j's goal at the first stage of the ordinary contracting game is to minimize the payoff of Principal q. Define Principal q's minmax value to be his lowest equilibrium payoff in this game. This minmax value is similar to the standard definition, except for the fact that principals can only punish Principal q in the contracting stage; each player behaves strategically in the subgame generated by the contract profile. In that subgame, others can only punish Principal q by playing an equilibrium which makes him worst off.

We characterize the set of subgame-implementable allocations in terms of the preferences of the agents and the principals. In particular, we show that an allocation is subgame-implementable only if it is *strongly incentive compatible*. To explain the notion of strong incentive compatibility, recall that an allocation in our model is a mapping from the agents' type spaces to action profiles

[3] That is, an ordinary contract is a mapping from the message profile of the agents to the subsets of the action space of the principal.

of the principals. Each coordinate of the allocation maps the vector of type profiles of the agents to the action space of a certain principal. Suppose that principals act simultaneously, each offering a direct mechanism to implement his coordinate of the allocation. An allocation is said to be strongly incentive compatible if truth-telling by all agents constitutes an equilibrium in the product of these direct mechanisms. That is, an allocation is strongly incentive compatible if no agent is able to increase his payoff by misreporting his type to the principals. This definition is stronger than the standard definition of incentive compatibility because, in our setting, agents may report different types to different principals.

The most conceptually challenging aspect of our folk theorem concerns these minmax values. Since contracts are contractible, one might imagine that the punishment inflicted on a deviating principal might depend on the actual deviation. If punishments could be made contingent on the deviator's contract, then one might suspect that a deviator could be pushed below his minmax value, perhaps even to his maxmin value. This argument turns out to be false. Despite the contractibility contracts, punishments can only depend on the deviator's identity, and not on his contract. In other words, when the principals punish a deviator they make use of ordinary contracts, and do not condition the punishment itself on the other principal's contract. This fact is due to an argument based on mathematical logic stated in Proposition 1.

Finally, in both ordinary and contractible contract settings, we identify equilibrium allocations which can only be implemented by contract profiles which do not pin down single actions for the principals in the last stage of the game. In this sense, equilibrium contracts are often incomplete. This is due to the existence of a trade-off between committing to a small set of actions and having flexibility at the last stage of the game. On the one hand, more commitment can increase ex ante payoffs. On the other hand, more flexibility can deter certain deviations. Indeed, a deviation might be more attractive if the deviator knows exactly what actions his opponents will take at the last stage of the game. There is another sense in which restricting attention to complete contracts results in a loss of generality. We come across allocations which can be supported as an equilibrium if contracts are required to be complete, but cannot be supported if contracts are allowed to be incomplete. This is because a principal might profitably deviate by offering an incomplete contract, but there might be no such a deviation in the form of a complete contract. These observations might provide new insights as to why contracts are often incomplete in the real world.

Literature Review

There is a sizeable applied literature on the theory of price relation agreements. As in our paper, these papers tend to conclude that price relation agreements facilitate tacit collusion. While we analyze general principal-agent models, this existing literature focuses only on price competition among firms. It is typically assumed that the contract of a firm consists of a posted price and a policy which maps the competitors' posted prices into prices. The actual price paid by a consumer for a firm's product is a function of the firm's posted price and the firm's policy evaluated at the others' posted prices, see, for example, Salop (1986), Png and Hirshleifer (1987), Belton (1987), Logan and Lutter (1989), Baye and Kovenock (1994), Chen (1995) and Zhang (1995). In other words, certain parts of the contracts are contractible, i.e. posted prices, but other parts of the contracts are not contractible, i.e. policies. It is unclear whether the conclusions of this literature are valid only taking into account these restrictions or whether they can be generalized to arbitrary contractible contracts. In addition, there are two drawbacks of these contract spaces. First, such contracts make sense only in the simplest price-setting contexts, and it is not obvious how they might be adapted to environments where the contractible objects are more complex than simple prices. Second, these contracts do not accommodate any communication between principals and agents, and hence cannot be used for screening. Indeed, the agents (buyers) in the models mentioned above are essentially non-strategic and possess no private information. In contrast, the contract space in our model is not context-specific, and can handle arbitrarily complex contractible decisions as well as adverse selection among agents.

In terms of results, our paper generalizes insights from this existing literature to arbitrary common-agency environments. Indeed, we show that the contractibility of contracts might lead to a softening of competition, and reduce welfare in an exclusive-contracting environment. In the context of price-setting, collusive contracts specify high prices if other principals also offer collusive contracts and trigger low prices if a principal deviates. This has a similar flavor to the *meet-the-competition clause*, which enables a firm to lower its price in response to undercutting by its competitor. The empirical literature on the meet-the-competition clause is ample and is consistent with our results. The seminal paper is by Hess and Gerstner (1991) which analyzes the competition between two large supermarkets, Winn Dixie and Food Lion, in North Carolina. The authors document that Winn Dixie's adoption of the meet-the-competition clause led to coordinated prices which ultimately reduced consumer welfare. Arbatskaya et al. (2004) analyze the relationship between prices and advertised price-matching promises of retailers across various industries. The authors confirm that price-matching promises typically soften competition. Arbatskaya et al. (2004) draws similar conclusions in the context of tire prices and price-match guarantees advertised in local newspapers.

Collusion among principals does not emerge as a unique equilibrium in our model. The contractibility of contracts does facilitate collusion, but also allows principals to behave competitively. In particular, a contract may specify a relatively low price irrespective of the pricing policies of competitors. An example of such a contract is a *price-beating guarantee* which strictly undercuts the prices of competitors. Not surprisingly, most empirical studies find that such guarantees do not lead to softening of the competition, see, for example, Arbatskaya et al. (2004, 2006) and Manez (2006). Whether or not price-relationship contracts across sellers should be prohibited is subject to ongoing debate. The empirical findings suggesting that a meet-the-competition clause softens competition but price-beating promises do not are interpreted as mixed evidence in favor of making these agreements illegal, see for example Aguzzoni et al. (2012). We adopt a different view; we propose that contracting on contracts should be disallowed because it fosters collusion among principals, even though it does not necessarily lead to it.

Our paper is also related to the theoretical literature on common agency. Our most important departure from this literature is the contractibility of contracts. Although the meet-the-competition example is frequently used as motivation, this literature usually assume that the contracts cannot be contracted upon directly, but only through the reports of the agents. In order for the agents to communicate their contracts to the principals, their message spaces must be at least as large as the space of contracts. Since the contracts are mappings from the message spaces, it is not straightforward to construct such a message space. Epstein and Peters (1999) show that there exists a universal message space that is rich enough for agents to communicate their private information as well as the contracts offered by the principals. They show that any equilibrium in a communication game with a large enough message space can be implemented as an equilibrium in the game with the universal message space.[4]

Peters (2001) and Martimort and Stole (2002) show that a version of the Taxation Principle holds for common agency games. That is, any equilibrium in any communication game can be implemented as an equilibrium in a game where the principals offer menus of ordinary contracts. An ordinary contract is one which maps reports of types to outcomes. The agent then selects items from the menu of each principal. One of the shortcomings of the literature is a lack of characterization of these allocations. Perhaps the main contribution of our paper to this literature is the full characterization of the equilibrium allocations.

Finally, our paper is also related to the literature on mutually dependent commitment devices, see for example Tennenholtz (2004), Kalai et al. (2010) and Peters and Szentes (2012). This literature considers two-stage games in which players submit commitment devices at the first stage, and play a normal form game at the second stage. A player's commitment device is a restriction of his action space as a function of the commitment devices of the other players. Various folk theorems have been proven in these situations. The equilibrium construction is usually based on a *self-referential* commitment device, similar to the concept in our paper. Tennenholtz (2004) analyzes complete information games and models the commitment device space as the set of Turing machines. Tennenholtz (2004) proves one direction of a pure-strategy folk theorem. That is, he shows the implementability of any outcome in which each player receives at least his minmax payoff. Kalai et al. (2010) characterizes mixed-strategy equilibria in complete information environments. Their main theorem states that any *correlated* outcome of the second-stage normal form game can be implemented by commitment devices in which all players' payoffs exceed their minmax payoffs.

Peters and Szentes (2012) departs from complete information environments and investigates Bayesian games with commitment devices. As in this paper, Peters and Szentes (2012) models commitment devices as definable functions, and also show that this space includes self-referential devices and that a player's payoff cannot be pushed below his minmax value. In contrast, the

[4] Calzolari and Pavan (2006) and Yamashita (2010) simplify the universal message space, which makes it possible to characterize equilibria in special cases.

problem of competitive screening does not arise in the model of Peters and Szentes (2012) because they do not have agents. The key trade-off in Peters and Szentes (2012) is related to the information content of the devices. On the one hand, a player benefits from offering type-contingent devices because different types prefer to commit to different actions. On the other hand, a player might be hurt at the second stage if he reveals too much information through his devices. Therefore, different types of players might prefer to offer the same device in order to disclose less information. A player's equilibrium devices balance these countervailing incentives and generate a partition of the player's type space. Two different types of players submit the same device if and only if they belong to the same partition element. Let us emphasize that no such trade-off is present in this paper, as the principals have no information to start with. The main result of Peters and Szentes (2012) states that the set of allocations implementable with their contracting game is the same as the set of allocations implementable with public message mechanisms. A public message mechanism is similar to a standard direct mechanism except that messages are publicly observable and non-participants can arbitrarily restrict their action spaces as a function of others' reports.

2 An Example

The goal of this section is to explain the space of contracts and provide an illustration of our approach in the context of an oligopoly example.

Suppose there are two firms (1 and 2) and a single consumer. Each firm can produce a particular good at no cost. The goods are close substitutes but not identical. The consumer has one of two equally likely types, A and B. The consumer's valuations for the goods are summarized in the following table:

	A	B
Firm 1	9	8.5
Firm 2	8.5	9

That is, if the consumer's type is A, he values Firm 1's good at 9 and Firm 2's good at 8.5. If his type is B, he values Firm 2's good at 9 and Firm 1's good at 8.5. His marginal value for a second good is zero. The action space of each firm consists of setting a price from the set $\{0, ..., 10\}$. The firms maximize their profits, and the agent wants to maximize his value for the good he purchases minus the price. Consider the following game. First, firms submit contracts simultaneously. A contract specifies a price as a function of the contract of the other firm and the message of the consumer. Contracts are publicly observable. Second, the consumer sends messages to each firm. Finally, the consumer decides which product to buy if any. If the firms were to set prices simultaneously without being able to contract on contracts, the market prices and the joint profit would be at most 2. If the firms could collude, they would maximize their joint profit by setting a price of 9 and the consumer would buy the good from Firm 1 if his type is A and from Firm 2 if his type is B. Next, we explain the contract space and show that it is possible to implement the collusive outcome with contracts.

We endow each market participant with a formal language. We require each contract offered by a firm and each message sent by the consumer to be a text written in this language, where a text is a finite string of symbols. A firm's contract gives *precise instructions* on how to determine the price as a function of the texts submitted by the firm and the consumer. Below, we construct a contract for each firm, c_1^* and c_2^*, which implements the allocation that maximizes the joint profit. These contracts will be cross-referential and take the following form

$$c_1^*(c_2, m) = \begin{cases} 9 & \text{if } c_2 = c_2^*, \text{ and } m = A \\ 10 & \text{if } c_2 = c_2^*, \text{ and } m \neq A \\ 0 & \text{otherwise,} \end{cases} \quad \text{and } c_2^*(c_1, m) = \begin{cases} 9 & \text{if } c_1 = c_1^*, \text{ and } m = B \\ 10 & \text{if } c_1 = c_1^*, \text{ and } m \neq B \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where c_1 and c_2 denote the contracts of Firms 1 and 2, respectively, and m is the consumer's message. The contract c_1^*, for example, specifies a price of 9 if Firm 2's contract is c_2^* and the consumer's report is A and a price of 10 if Firm 2's contract is c_2^* and the consumer's report is not A. If the contract of Firm 2 is different from c_2^* then the contract c_1^* forces Firm 1 to set a price of zero. If the firms offer c_1^* and c_2^*, the best response of the consumer is to report her type truthfully. Also note that if Firm i offers c_i^*, the best response of Firm j $(j \neq i)$ is to offer c_j^* because any other contract would trigger a price of zero by Firm i. So the contract profile (c_1^*, c_2^*) indeed implements an allocation which maximizes the joint profit. The problem is that c_1^* explicitly depends on c_2^* and c_2^* explicitly depends on c_1^*. So, the text corresponding to c_1^* has to include a description of the text describing c_2^*, which, in turn, has to include the description of c_1^* itself. Below, we explain how to construct these texts.

We take advantage of the fact that languages can be coded. That is, there exists a bijection from the set of finite texts into the set of integers, so each text can be coded by a unique integer. One such mapping is called the *Gödel Coding*. So, to any text which describes a mapping from texts to prices there is a corresponding text which describes a mapping from codes of texts into prices. Since all the codes are integers, this latter text is a description of an arithmetic mapping. An arithmetic function which can be described in a formal language is called definable. This set is formally defined in the next section. Since the Gödel Coding can also be described using a text, we can identify the space of contracts with the set of definable functions from $\mathbb{N}^2 \to \{0, ..., 10\}$, where the first argument of these functions is the code for the other firm's contract and the second argument is the code for the message of the consumer. The range for these functions is the set of prices.[5] In what follows, we use Gödel Coding to construct texts corresponding to c_1^* and c_2^* in (1).

Let $[\varphi]$ denote the Gödel code of the text φ. Consider the following contract for Firm 1:

$$c_1^{n_2}([c_2], [m]) = \begin{cases} 9 & \text{if } [c_2] = n_2, \text{ and } [m] = [A] \\ 10 & \text{if } [c_2] = n_2, \text{ and } [m] \neq [A] \\ 0 & \text{otherwise.} \end{cases}$$

[5] In general, the actions of a principal do not correspond to integers. Then the range of these functions are the codes of the actions.

This contract says that if the Gödel code of Firm 2's contract is n_2 and the consumer reports type A, then Firm 1 will set the price at 9. If the Gödel code of Firm 2's contract is n_2 but the consumer does not report type A, then the price will be 10. Otherwise, the price will be zero. Similarly, define Firm 2's contract as follows:

$$c_2^{n_1}\left([c_1],[m]\right)=\begin{cases} 9 & \text{if } [c_1]=n_1, \text{ and } [m]=[B] \\ 10 & \text{if } [c_1]=n_1, \text{ and } [m]\neq[B] \\ 0 & \text{otherwise.} \end{cases}$$

Notice that if $[c_1^{n_2}]=n_1$ and $[c_2^{n_1}]=n_2$ then these contracts correspond to c_1^* and c_2^*. Therefore, we have reduced the problem of constructing cross-referential contracts to finding a solution to the fixed-point problem $(n_1,n_2)=\left([c_1^{n_2}],[c_2^{n_1}]\right)$.

Before we proceed, we introduce two pieces of notation. First, the function $<.>$ is the Gödel coding inverse operation. That is, $<n>$ is the text whose Gödel code is n. Second, we shall make use of *free variables* to express statements such as "$x>y$" in texts. Integers can be substituted for the free variables in order to make statements about these integers. If ϕ is a text, then $\phi^{(n_1,n_2)}$ denotes the same text as ϕ, except that if ϕ contained the free variables x or y, then the value of the free variable x is set to be n_1 and the value of the free variable y is set to be n_2. For example, if $\phi\equiv$"$x>y$" then $\phi^{(3,2)}\equiv$"$3>2$." Now, consider the following two texts:

$$c_1^{x,y}\left([c_2],[m]\right)=\begin{cases} 9 & \text{if } [c_2]=\left[<y>^{(x,y)}\right], \text{ and } [m]=[A] \\ 10 & \text{if } [c_2]=\left[<y>^{(x,y)}\right], \text{ and } [m]\neq[A] \\ 0 & \text{otherwise,} \end{cases}$$

$$c_2^{x,y}\left([c_1],[m]\right)=\begin{cases} 9 & \text{if } [c_1]=\left[<x>^{(x,y)}\right], \text{ and } [m]=[B] \\ 10 & \text{if } [c_1]=\left[<x>^{(x,y)}\right], \text{ and } [m]\neq[B] \\ 0 & \text{otherwise.} \end{cases}$$

These texts are not contracts, because they contain free variables. However, they become contracts when we evaluate these free variables at any pair of integers. Let γ_1 and γ_2 denote the Gödel codes of these two texts respectively. Then

$$c_1^{\gamma_1,\gamma_2}\left([c_2],[m]\right)=\begin{cases} 9 & \text{if } [c_2]=\left[<\gamma_2>^{(\gamma_1,\gamma_2)}\right], \text{ and } [m]=[A] \\ 10 & \text{if } [c_2]=\left[<\gamma_2>^{(\gamma_1,\gamma_2)}\right], \text{ and } [m]\neq[A] \\ 0 & \text{otherwise,} \end{cases} \tag{2}$$

and

$$c_2^{\gamma_1,\gamma_2}\left([c_1],[m]\right)=\begin{cases} 9 & \text{if } [c_1]=\left[<\gamma_1>^{(\gamma_1,\gamma_2)}\right], \text{ and } [m]=[B] \\ 10 & \text{if } [c_1]=\left[<\gamma_1>^{(\gamma_1,\gamma_2)}\right], \text{ and } [m]\neq[B] \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Recall that γ_1 is the Gödel of $c_1^{x,y}$, so $<\gamma_1>^{(\gamma_1,\gamma_2)}$ is $c_1^{\gamma_1,\gamma_2}$. Similarly, $<\gamma_2>^{(\gamma_1,\gamma_2)}$ is just $c_2^{\gamma_1,\gamma_2}$. Therefore, one can replace $<\gamma_1>^{(\gamma_1,\gamma_2)}$ and $<\gamma_2>^{(\gamma_1,\gamma_2)}$ with $c_1^{\gamma_1,\gamma_2}$ and $c_2^{\gamma_1,\gamma_2}$ in (2) and (3) and conclude that $c_1^{\gamma_1,\gamma_2}$ and $c_2^{\gamma_1,\gamma_2}$ are indeed the cross-referential contracts corresponding to (1). What's more, each principal's contract is now well-defined: each contract gives precise instructions as to how the consumer's message and the exact content of the other principal's contract together dictate the price that a principal will offer.

3 The Model

3.1 The Physical Environment

There are n principals and k agents. Each principal has a finite action space. The set of actions available to Principal j is $A_j = \times_{i=1}^{k} A_j^i$, where A_j^i denotes the set of actions of Principal j which affects the payoff of Agent i. Let A and A^i denote $\times_{j=1}^{n} A_j$ and $\times_{j=1}^{n} A_j^i$, respectively. The finite type space of Agent i is T^i, and T denotes $\times_{i=1}^{k} T^i$. The joint distribution of types is common knowledge. The payoff to Principal j is given by $u_j : T \times A \to \mathbb{R}$. The payoff to Agent i is $v_i : T \times A^i \to \mathbb{R}$. Principals and agents all maximize expected utility.

3.2 The Language and the Gödel Coding

We consider a formal language that is sufficiently rich to allow its user to state any arithmetic proposition. This implies that one can express, for example, that there exist Pythagorean triples:

$$\exists x, y, z \left[(n \geq 3)\, (x \neq 0) \vee (y \neq 0) \vee (z \neq 0) \vee (x^2 + y^2 = z^2) \right].$$

In addition, statements that involve any finite number of free variables can be expressed. For example,"$x < 4$" is a sentence in our language and the symbol x is a free variable in the statement. One can substitute any integer into x and then the predicate is either true or false. This particular sentence is true if $x = 1, 2, 3$ and false otherwise.

Definition 1 *The function $f : \mathbb{N}^k \to 2^{\mathbb{N}}$ is said to be* definable *if there exists a first-order arithmetic statement, ϕ, in $k+1$ free variables such that $b \in f(a_1, ..., a_k)$ if and only if $\phi(a_1, ..., a_k, b)$ is true.*

We provide the formal definition for a first-order arithmetic statement in the Appendix. The reader should keep in mind that a correspondence is definable if it can be *explained* in a language. To better understand the definition, consider the following correspondence: $f(n) = \{n, n+1\}$ for all $n \in \mathbb{N}$. In order to show that this correspondence is definable, we must construct the statement required by the previous definition. Let

$$\phi(x, y) \equiv (y = x) \vee (y = x + 1).$$

Notice that for any pair of integers, a and b, $\phi(a, b)$ is true if and only if b is either a or $a + 1$. Therefore, the predicate ϕ indeed defines f.

Let $\mathfrak{L}$ be the set of all sentences in our formal language. Each of its elements is a finite string of symbols. It is well known that one can construct a one-to-one function mapping $\mathfrak{L} \to \mathbb{N}$. Let $[\varphi]$ be the value this function takes at $\varphi \in \mathfrak{L}$, and call it the Gödel Code of the text φ.

3.3 The contracting game

Each principal offers a contract to each agent. The set of feasible contracts is the set of definable mappings from $\mathbb{N}^{nk} \times \mathbb{N} \to 2^{\mathbb{N}}$. The first nk arguments are the Gödel codes for the principals' contracts. The last argument is the code for the message sent by the agent to whom the contract is offered. We denote the set of contracts Principal j can offer to Agent i by C_j^i, and set $C_j = \times_{i=1}^k C_j^i$, $C = \times_{j=1}^n C_j$. The timing of the game is as follows. Principals simultaneously submit contracts $(c_1, \ldots, c_n) \in C$. These contracts are publicly observable. Then, agents send messages to the principals privately. Let m_j^i denote the message sent by Agent i to Principal j. Finally, principals simultaneously take actions chosen from the subsets of their action spaces determined by the contracts and messages. That is, Principal j can take action $a_j = \left(a_j^1, ..., a_j^k\right) \in A_j$ only if for all $i = 1, ..., k$

$$\left[a_j^i\right] \in c_j^i \left([c_1], ..., [c_n], \left[m_j^i\right]\right),$$

where $[c_q]$ denotes $\left(\left[c_q^1\right], ..., \left[c_q^k\right]\right)$ for all $q = 1, ..., n$. In the name of transparency, we will abuse notation and replace codes with actions, writing $c_j^i : \mathbb{N}^{nk} \times \mathbb{N} \to 2^{A_j^i} \setminus \{\emptyset\}$ while still thinking of c_j^i as a definable function.

These contracts can be implemented by cross-referential menus. The items on a menu of Principal j offered to Agent i are subsets of A_j^i. The menu corresponding to the contract c_j^i is $\left\{c_j^i \left([c_1], ..., [c_n], \left[m_j^i\right]\right) : \ m_j^i \in \mathbb{N}\right\}$ given c. Agent i's report can be interpreted as choice from this menu.

We restrict attention to pure-strategy perfect Bayesian equilibria (PBE). That is, the principals and agents are required to play a Bayesian equilibrium in every subgame generated by a contract profile.[6] The main result of this paper does not actually depend on the equilibrium concept, so long as players play *some* equilibrium in the subgames generated by the first-stage contracts. In particular, the set of sequential equilibria would be characterized by essentially the same constraints.

We also point out that the existence of an equilibrium is only guaranteed if mixed strategies are allowed at the second and third stages. The restriction to pure strategies is purely for notational convenience. Allowing for mixed strategies has no substantive consequence on our analysis.

[6] In order to guarantee that these subgames exist, one should describe the game such that the types of the agents are determined only after the contracts are offered by the principals. This way of modeling the game has no strategic implications but makes our terminology precise.

4 Equilibrium Characterization

We seek to characterize the set of allocations which can be implemented as PBE of the contracting game. A deterministic allocation is a mapping from the type profile of the agents to the action profiles of the principals. Our strategy is to first analyze equilibria in games where contracts are observable but not contractible. We call these games *ordinary contracting games*. The analysis of these games leads to a full characterization of the contractible contracting games. However, these games are interesting in their own right. In fact, these are the *communication games* analyzed in the common agency literature. In addition, we aim to identify environments where the ability to contract on contracts can lead to inefficiency. In order to do so, we have to characterize the set of equilibria in the ordinary contracting games.

4.1 Ordinary Contracting Games

The set of *ordinary contracts* is the set of definable mappings from $\mathbb{N} \to 2^{\mathbb{N}}$. The domain of these functions are the Gödel codes of the messages sent by the agent to whom the contract is offered. Let D_j^i denote the set of contracts Principal j can offer to Agent i, and let $D_j = \times_{i=1}^k D_j^i$ and $D = \times_{j=1}^n D_j$. The timing of the ordinary contracting game is as follows. Principals simultaneously select contracts $(d_1, \ldots, d_n) \in D$. These contracts are publicly observable. Then, agents send messages to the principals privately, $\{m^1, ..., m^k\} \subset \mathbb{N}^k$. Finally, principals take actions simultaneously, such that Principal j can take action $a_j = (a_j^1, ..., a_j^k) \in A_j$ if

$$[a_j^i] \in d_j^i\left([m_j^i]\right)$$

for all $i = 1, ..., k$. Again, for simplicity we use actions of the principals instead of their codes and write $d_j^i : \mathbb{N} \to 2^{A_j}$ while still thinking of d_j^i as a definable function. We restrict attention to pure-strategy PBE of this game.

We characterize the equilibria in these games by describing the best-response constraints of the principals and the agents. Notice that when an agent decides what messages to send to the principals, he knows his type and already observed the contract profile of the principals. Hence, the messages of the agents are functions of these two objects. Let $\beta^i : T^i \times D \to \mathcal{L}^n$ denote the strategy of Agent i, and let β_j^i denote the jth coordinate of β^i, that is, the message sent to Principal j by Agents i. Let β_j denote the messages received by Principal j, that is, $\left(\beta_j^1, ..., \beta_j^k\right)$. Principal j's action at the last stage of the game can depend on both the first-stage contract profile and the messages sent to him by the agents. Let $\alpha_j = \left(\alpha_j^1, ..., \alpha_j^k\right)$, $\alpha_j^i : \mathcal{L}^k \times D \to A_j^i$ for all i, denote the strategy of the principals at the last stage. Since Principal j's action must be consistent with his contract, $\alpha_j^i(m_j, d) \in d_j^i\left([m_j^i]\right)$ must hold for all i, $m_j \in \mathcal{L}^k$, and for all $d = (d_j, d_{-j}) \in D$. As usual, α_{-j} denotes the action profile of principals other than Principal j, and β^{-i} denotes the message profile of agents other than Agent i.

In what follows, we define PBE in terms of three sets of constraints. The first constraint guarantees that each principal takes an action at the last stage which maximizes his payoff. For all j, $d \in D$:

$$\alpha_j\left(m_j, d\right) \in \arg \max_{a_j \in d_j\left(\left[m_j\right]\right)} E_t\left[u_j\left(t, a_j, \alpha_{-j}\right) \mid d, m_j, \beta, \alpha_{-j}\right], \tag{4}$$

for all $m_j \in \mathcal{L}^k$ and $d \in D$. The expectations are formed according to Bayes Rule if the message profile sent by the agents, m_j, is consistent with their equilibrium behavior. However, PBE imposes no restriction on the belief of Principal j if m_j is off the equilibrium path.[7]

The second constraint ensures that each agent maximizes his payoff by his message in every subgame generated by a contract profile. For all i, $t^i \in T^i$, and $d \in D$,

$$\beta^i\left(t^i, d\right) \in \arg \max_{m^i \in \mathcal{L}^n} E_{t^{-i}}\left[v_i\left(t, \alpha^i\left(\left(m^i, \beta^{-i}\right), d\right)\right) \mid d, t^i\right]. \tag{5}$$

The last constraint guarantees that no principal wants to deviate from his equilibrium contract in the first stage of the game. Let $\left(d_1^*, ..., d_n^*\right) = d^*$ denote the equilibrium contract profile. Then, for all j:

$$d_j^* \in \arg \max_{d_j \in D_j} E_t\left(u_j\left(t, \alpha\left(\beta, \left(d_j, d_{-j}^*\right)\right)\right)\right). \tag{6}$$

We are ready to define PBE as follows:

Definition 2 *The strategy profile (d^*, β, α) constitutes a PBE in the Ordinary Contracting Game if and only if (4), (5), and (6) are satisfied.*

It turns out to be useful to define the set of those allocations that can be implemented in a subgame generated by *some* ordinary contract profile. To this end, let σ^d denote the set of those (α, β) pairs for which both (4) and (5) are satisfied. Then the set of allocations that can be implemented in some subgame is defined as follows:

$$\mathcal{A} = \left\{g : T \to A : \exists d \in D,\ \exists (\alpha, \beta) \in \sigma^d \text{ s.t. } g(t) = \alpha(\beta(t, d), d)\right\}.$$

We refer to $\mathcal{A}$ as the set of *subgame-imlementable* allocations. Next, we characterize this set in terms of the preferences of the agents and the principals.

First, we fix an allocation g and explore the implications of $g \in \mathcal{A}$ for the preferences of the agents. Let $(\alpha, \beta) \in \sigma^d$ for some $d \in D$ such that $g(t) \equiv \alpha(\beta(t, d), d)$. Consider Agent i with type t^i and fix an arbitrary vector $\left(t_1^i, ..., t_n^i\right) \in \left(T^i\right)^n$. Then, by (5), Agent i is better off sending the message profile $\beta^i\left(t^i, d\right)$ as opposed to $\beta_j^i\left(t_j^i, d\right)$ to each Principal j. This implies

$$E_{t_{-i}}\left[v^i\left(t, \left(g^i\left(t^i, t^{-i}\right)\right)\right) \mid t^i\right] \geq E_{t_{-i}}\left[v^i\left(t, \left(g_1^i\left(t_1^i, t^{-i}\right), ..., g_n^i\left(t_n^i, t^{-i}\right)\right)\right) \mid t^i\right]. \tag{7}$$

Indeed, the left-hand side of this inequality is the expected payoff of Agent i conditional on t^i in the subgame generated by d and given (α, β). The right-hand side is the expected payoff of Agent i conditional on t^i if he deviates and sends message $\beta_j^i\left(t_j^i, d\right)$ to Principal j instead of $\beta_j^i\left(t^i, d\right)$. The inequality (7) motivates the following

[7] A stronger equilibrium refinement concept imposes restrictions on the beliefs according to which the expectations are formed in (4), but has no other impact on our characterization result.

Definition 3 *Let* $g_j^i : T \to A_j^i$ *for all* $j = 1, ..., n,\ i = 1, ..., k$ *and let* $g^i = \left(g_1^i, ..., g_n^i\right)$*. Then the allocation* $g = \left(g^1, ..., g^k\right)$ *is called* strongly incentive compatible *if for all* $i \in \{1, ..., k\}$*,* $t^i \in T^i$*, and* $\left(t_1^i, ..., t_n^i\right) \in \left(T^i\right)^n$ *the inequality (7) is satisfied.*

This definition is simply the standard notion of incentive compatibility extended to a multi-principal setting. Indeed, this definition would coincide with the standard definition if the inequality (7) were required to hold only for those type vectors, $\left(t_1^i, ...t_n^i\right) \in \left(T^i\right)^n$, where $t_1^i = t_2^i = ... = t_n^i$. Such a constraint would require that no agent be able to benefit from mimicking another of his type. In our multi-principal model, however, we must take more complex deviations into account. In particular, the messages of the agents are private, and therefore, an agent may report different types to different principals. Of course, any strongly incentive compatible allocation will also be incentive compatible. The following example shows that the converse is not true.

Example 1. Suppose that $n = 2$, $k = 1$, and $A_1 = A_2 = \{a_1, a_2\}$. The agent has two equally likely types, $T = \{1, 2\}$. The payoffs to the agent are described by the following matrix:

$$\begin{array}{c|c|c|} & a_1 & a_2 \\ \hline a_1 & 0 & 1 \\ \hline a_2 & 1 & 0 \\ \hline \end{array}.$$

The allocation g, defined by $g(t) = (a_t, a_t)$ for $t = 1, 2$, is obviously incentive compatible but not strongly incentive compatible.

Next, we turn our attention to the principals. In the subgame generated by the contract profile d, Agents with different types might send the same message to Principal j. The reporting strategy of Agent i generates a partition on T^i denoted by $\tau_j^i : T^i \to 2^{T^i} \setminus \{\emptyset\}$.[8] Let τ_j denote $\times_{i=1}^n \tau_j^i$. After receiving the messages, Principal j learns only $\tau_j(t)$ but not t. Since the action of Principal j can depend only on information he knows, the strategy profile (α, β) implements the allocation g only if the function $g_j(t)$ is measurable with respect to τ_j, that is, $g_j(t) = g_j(t')$ whenever $\tau_j(t) = \tau_j(t')$. In addition, the action of Principal j must be consistent with his contract, that is, $\alpha_j^i\left(m_j^i, d\right) \in d_j^i\left(\left[m_j^i\right]\right)$. As a consequence, the set $d_j^i\left(\left[\beta_j^i(t_i, d)\right]\right)$ must contain $g_j^i\left(t^i, t^{-i}\right)$ for all $t^{-i} \in T^{-i}$. Therefore, an implication of (4) is that

$$g_j(t) \in \arg \max_{a_j^i \in \left\{g_j^i\left(t^i, t^{-i}\right) \,:\, t^{-i} \in T^{-i}\right\}} E_t\left[u_j\left(t, a_j, g_{-j}(t)\right) \mid \tau_j(t)\right]. \tag{8}$$

To summarize, we have argued that if $g \in \mathcal{A}$ then g is strongly incentive compatible, and there is a partition of the type space for which (8) holds. Next, we show that the converse is also true. The following lemma fully characterizes the set of subgame-implementable allocations.

Lemma 1 *The allocation* $g : T \to A$ *is an element of* $\mathcal{A}$ *if and only if*

(i) g *is strongly incentive compatible and*

(ii) there exists a partition, $\tau_j^i : T^i \to 2^{T^i} \setminus \{\emptyset\}$ *for all* (i, j) *such that* g_j *is* τ_j*-measurable and (8) is satisfied for all* $j = 1, ..., n$*.*

[8] That is, $\tau_j^i\left(t_1^i\right) = \tau_j^i\left(t_2^i\right)$ if and only if $\beta_j^i\left(t_1^i, d\right) = \beta_j^i\left(t_2^i, d\right)$.

Proof. The "only if" part of the proof is already established in the text. To prove the "if" part, suppose that (i) and (ii) are satisfied. Consider the following contract offered to Agent i by Principal j:

$$d_j^i\left(\left[m_j^i\right]\right)=\begin{cases}\left\{g_j^i\left(t^i,t^{-i}\right)\ :\ t^{-i}\in T^{-i}\right\} & \text{if } m_j^i=\tau_j^i\left(t^i\right)^{9} \text{ and}\\ \underline{a}_j^i & \text{otherwise,}\end{cases}$$

where $\underline{a}_j^i$ is an arbitrary element of A_j^i. By (ii), the function g_j is measurable with respect to τ_j, so this contract is well-defined. Since the allocation g is strongly incentive compatible, truth-telling by the agents constitutes an equilibrium in the subgame. (That is, $m_i^j\left(t^i,d\right)=\tau_j^i\left(t^i\right)$ for all i, t^i and j is an equilibrium.) Finally, by (8), Principal j optimally chooses action $g_j\left(t\right)$ if the type profile of the agents is t. This equilibrium obviously implements g. ■

In the case with a single agent, the contracts constructed in the proof of the previous lemma determine single actions for the principals as a function of the agent's message profile. In the subgames generated by these contracts, the principals do not make any decisions, and hence, part (ii) is always satisfied. Therefore, we claim the following

Remark 1 *Suppose that $k=1$. Then the allocation $g:T\rightarrow A$ is an element of $\mathcal{A}$ if and only if g is strongly incentive compatible.*

We further investigate the properties of equilibria of the ordinary contracting games in Section 6. Next, however, we use Lemma 1 to characterize the set of equilibria in contractible contracting games.

4.2 Contractible Contracting Games

This section is devoted to the characterization of the equilibria in the contractible contracting game. We prove a folk theorem and show that an allocation is implementable if and only if it subgame-implementable and the payoff of each principal is larger than his minmax value, to be defined later. To see that the allocation must be subgame-implementable, we first argue that any contract profile generates an ordinary contract profile. To this end, suppose that $\left(c_1^*,...,c_n^*\right)$ is an equilibrium contract profile. For each (j,i), define $d_j^{i*}\in D_j^i$, such that $d_j^{i*}\left(l\right)=c_j^{i*}\left(\left[c_1^*\right],...,\left[c_n^*\right],l\right)$ for all $l\in\mathbb{N}$ and let d_j^* denote $\left(d_j^{1*},...,d_j^{k*}\right)$. Notice that $d^*=\left(d_1^*,...,d_n^*\right)$ is an ordinary contract profile, and the subgame generated by c^* in the contractible contracting game is the same as the subgame generated by d^* in the ordinary contracting game. Since players are required to play an equilibrium in the subgame generated by the first-stage contract profile, we can conclude that any allocation that can be implemented as a PBE in the contractible contracting game must belong subgame-implementable.

The difficult part of the theorem is to pin down the minmax values of the principals. The minmax value of Principal j is the lowest possible value that he can get in the ordinary contracting

[9] To be more precise, m_j^i is a text describing $\tau_j^i\left(t^i\right)$.

game if the goal of the other principals at the first stage of the game is to minimize his payoff. Formally, we shall prove that the minmax value of Principal j, $\underline{u_j}$, is:

$$\underline{u_j} = \min_{d_{-j} \in D_{-j}} \max_{d_j \in D_j} \min_{(\alpha,\beta) \in \sigma^{(d_j, d_{-j})}} E_t \left(u_j \left(t, \alpha \left(\beta \right) \right) \mid \left(d_j, d_{-j} \right) \right). \tag{9}$$

The meaning of this expression can be explained as follows. All the principals other than Principal j offer ordinary contracts at the first stage of the game in order to minimize the payoff of Principal j. Principal j also offers an ordinary contract which is a best response to the contracts of the others. These contracts generate a subgame in which there can be multiple equilibria. In this subgame, the principals and agents play an equilibrium which is the worst one for Principal j.

The fact that Principal j can only be punished by playing the worst equilibrium in the subgame is obvious because PBE requires the players to play an equilibrium in any subgame generated by a contract profile. The nontrivial part of our main result is the rest of the definition of $\underline{u_j}$. As we explained at the beginning of this section, the equilibrium contracts and a first-stage deviation of Principal j determines an ordinary contract profile. The formula in (9) essentially says that the ordinary contract profile of the principals other than Principal j does not depend on the deviation of Principal j, and hence, Principal j can best-respond to it. Since contracts are contractible, the ordinary contract profile of the principals other than Principal j can depend on the deviation of Principal j. Therefore, one might conjecture that the principals might be able to push Principal j's value below $\underline{u_j}$. For example, if Principal j would be restricted to offer ordinary contracts then the others could always offer contracts which are contingent on the ordinary contract of Principal j. Being able to offer these contingent contracts, is similar to being able to move after observing Principal j's contract, and hence, his lowest value would be

$$\max_{d_j \in D} \min_{d_{-j} \in D_{-j}} \min_{(\alpha,\beta) \in \sigma^{(d_j, d_{-j})}} E_t \left(u_j \left(t, \alpha \left(\beta \right) \right) : \left(d_j, d_{-j} \right) \right).$$

Of course, Principal j is not restricted to offer ordinary contracts, and his contract can be contingent on the contracts offered by the other principals, which are contingent on his contract etc. In fact, because of this infinite regress problem, it is not even clear that the lowest value of Principal j is well-defined.

Nevertheless, we show that this value is well-defined and, interestingly, the most severe punishment inflicted on Principal j can be assumed to be invariant to his deviation. To be more specific, Proposition 1 shows that no matter what the contract profile of the principals is, there always exists an ordinary contract profile $d_{-j} \in D_{-j}$, such that for all $d_j \in D_j$, there is a way for Principal j to write a contract so that the generated ordinary contract profile is (d_j, d_{-j}). But then it is without the loss of generality to assume that the principals use the ordinary contract profile d_{-j} to punish Principal j.

We are ready to state our main result formally.

Theorem 1 *An allocation $g : T \to A$ is implementable as an equilibrium in the contractible contracting game if and only if (i) g is subgame-implementable, and (ii) for all $j \in \{1, ..., n\}$*

$$E_t u_j (t, g(t)) \geq \underline{u_j}.$$

We break the proof of the theorem into two parts. The "if" part is based on the same arguments as the ones used in the example of Section 2. We shall construct cross-referential contracts which support the desired allocation. Essentially, the contract of Principal j (for all j) specifies target codes, k for each of the other principals. If the Gödel codes of the contracts of Principal q are the same as his target codes for all q, then Principal j *cooperates*. If Principal q deviates, and the codes of his contracts are different from his target codes, the contract of Principal j prescribes an ordinary contract which is used to minmax Principal q. The set of equilibrium contracts are cross-referential because the Gödel codes of Principal j's contracts, which we have just described, are exactly the same as his target codes specified in the contracts of all the other principals.

Recall two pieces of notation from the introduction. First, if $l \in \mathbb{N}$ then $< l >$ denotes the text whose Gödel code is l. That is, $[< l >] = l$. Second, for any text φ and $(l_1, ..., l_n)$, let $\varphi^{(l_1,...,l_n)}$ denote the text where if the letter x_q stands for a free variable in φ then x_q is substituted for l_q in φ for $q = 1, ..., n$. For example, if φ is "$x_1 < x_2$", $l_1 = 1$, and $l_2 = 2$ then $\varphi^{(l_1,l_2)}$ is $1 < 2$.[10]

Consider now the following text in n free variable: $< x_q >^{(x_1,...,x_n)}$, where $q \leq n$. Since the Gödel coding is a bijection, $< l_q >$ is a text for each $l_q \in \mathbb{N}$. Since $\varphi^{(l_1,...,l_n)}$ is defined for all φ and $(l_1, ..., l_n) \in \mathbb{N}^n$, $< l_q >^{(l_1,...,l_n)}$ is a text for all $(l_1, ..., l_n) \in \mathbb{N}^n$. It is a well-known result in Mathematical Logic that if $f (l_1, ..., l_n) = \left[< l_q >^{(l_1,...,l_n)}\right]$, then f is a definable function.

Proof of the "if" part of Therorem 1. Since the allocation g is in $\mathcal{A}$ there exists an ordinary contract profile $d^* = (d_1^*, ..., d_n^*)$, a strategy profile of the agents, $\beta^* = \left(\beta^{*1}, ..., \beta^{*k}\right)$, a third-stage strategy profile of the principals, $\alpha^* = (\alpha_1^*, ..., \alpha_n^*)$, such that $g(t) = \alpha^* (\beta^* (t, d^*), d^*)$ and both (4) and (5) are satisfied, that is, $(\alpha^*, \beta^*) \in \sigma^{d^*}$. In addition, let $d_{j,q}$ denote the contracts of Principal j which he uses to minmax Principal q. That is, the contract profile $d_{-q,q}$ solves

$$\min_{d_{-q} \in D_{-q}} \max_{d_q \in D_q} \min_{(\alpha,\beta) \in \sigma^{(d_q,d_{-q})}} E_t (u_q (t, \alpha(\beta)) : (d_q, d_{-q})). \tag{10}$$

Let $x_m = \left(x_m^1, ..., x_m^k\right)$ a vector of free variables for all $m = 1, ..., n$. Consider the following text of Principal j, $c_j^{i,x_1,...,x_n}$, in nk free variables:

$$c_j^{i,x_1,...,x_n} \left(([c_l])_{l=1}^n, [m_j^i]\right) =$$

$$\begin{cases} d_j^{*i} \left([m_j^i]\right) & \text{if } \left|\left\{l : \exists i \text{ s.t. } \left[< x_l^i >^{(x_1,...,x_n)}\right] \neq \left[c_l^i\right] \right\}\right| \neq 1, \\ d_{j,q}^i \left([m_j^i]\right) & \text{if } \left\{l : \exists i \text{ s.t. } \left[< x_l^i >^{(x_1,...,x_n)}\right] \neq \left[c_l^i\right] \right\} = \{q\}, \end{cases} \tag{11}$$

[10] Of course, it is possible that the text φ does not contain some of the symbols $\{x_1, ..., x_n\}$. In that case, there is no substitution for the missing letters in $\varphi^{(l_1,...,l_n)}$. For example, if φ is "$x_2 > 2$", then $\varphi^{(3,4)}$ is "$4 > 2$", because x_1 does not appear in φ.

for all $\left[m_j^i\right] \in \mathbb{N}$. This expression (11) is not a contract, but rather a contract with free variables. However, $c_j^{x_1,\ldots,x_n}$ would become a contract if the free variables $(x_1, ..., x_n)$ are replaced by integers. Each of these contracts with free variables has a Gödel code, so let $\gamma_j^i = \left[c_j^{i,x_1,\ldots,x_n}\right]$ and $\gamma_j = \left(\gamma_j^1, ..., \gamma_j^k\right)$. The functions $\left\{c_j^{i,\gamma_1,\ldots,\gamma_n}\right\}_{i,j}$ have no free variables, so they constitute a set of contracts. Notice that

$$
\begin{aligned}
& c_j^{i,\gamma_1,\ldots,\gamma_n}\left(\left([c_l]\right)_{l=1}^n, \left[m_j^i\right]\right) \\
= & \begin{cases} d_j^{*i}\left(\left[m_j^i\right]\right) & \text{if } \left|\left\{l : \exists i \text{ s.t. } \left[<\gamma_l^i>^{(\gamma_1,\ldots,\gamma_m)}\right] \neq \left[c_l^i\right]\right\}\right| \neq 1, \\ d_{j,q}^i\left(\left[m_j^i\right]\right) & \text{if } \left\{l : \exists i \text{ s.t. } \left[<\gamma_l^i>^{(\gamma_1,\ldots,\gamma_m)}\right] \neq \left[c_l^i\right]\right\} = \{q\}. \end{cases}
\end{aligned} \tag{12}
$$

The contract $c_j^{i,\gamma_1,\ldots,\gamma_n}$ is definable because d_j^{*i}, $d_{j,q}^i$ and $f(l_1, ..., l_n) = \left[<l_q>^{(l_1,\ldots,l_n)}\right]$ are all definable. Observe what happens when Principal q offers contract $c_q^{i,\gamma_1,\ldots,\gamma_n}$ for all q, i. Principal j needs to check whether the Gödel code of $<\gamma_q^i>^{(\gamma_1,\ldots,\gamma_m)}$ is equal to the Gödel code of $c_q^{i,\gamma_1,\ldots,\gamma_n}$. The integer γ_q^i is the Gödel code of the contract with free variables $c_q^{i,x_1,\ldots,x_n}$. Principal j's contract says to take this contract with the free variables, fix the free variables at $\gamma_1, ..., \gamma_n$ (which gives the contract $c_q^{i,\gamma_1,\ldots,\gamma_n}$), then evaluate its Gödel code. This is what is to be compared with the Gödel code of the contract offered by Principal q to Agent i. Of course, if Principal q offers $c_q^{i,\gamma_1,\ldots,\gamma_n}$ to Agent i these are the same. In fact, if Principal q offers $c_q^{i,\gamma_1,\ldots,\gamma_n}$ for all (q, i) then Principal j ends up with the ordinary contract d_j^{*i} according to the first line of (12). Therefore, if Principal j offers contract $c_j^{i,\gamma_1,\ldots,\gamma_n}$ for all (j, i) then the resulting subgame is generated by the ordinary contract profile d^*. Define the strategies of the agents and the principals as $\beta^*(\cdot, d^*)$ and $\alpha^*(\cdot, d^*)$. These strategies obviously support the allocation g. It remains to specify the strategies of the players off the equilibrium path and show that no player can profitably deviate.

Next we define the second-stage strategies of the agents and the third-stage strategies of the principals off the equilibrium path. It is enough to define these strategies in subgames which result from a deviation of a single principal. Suppose that Principal q offers a contract c_q^i instead of $c_q^{i,\gamma_1,\ldots,\gamma_n}$ to Agent i. Let $d_q(\cdot)$ denote $c_q\left([c_q], \left(\left[c_j^{\gamma_1,\ldots,\gamma_n}\right]\right)_{j\neq q}, \cdot\right)$. As a result of this deviation, according to the second line of (12), Principal j will end up with the ordinary contracts $d_{j,q}$ for all $j \neq q$. Therefore, the subgame resulting from the deviation of Principal q is generated by the ordinary contract profile $d = (d_q, d_{-q,q})$. Define the strategies of the agents and the principals, $\alpha(d)$ and $\beta(d)$, so that the expected payoff of Principal q is minimized. That is, $(\alpha(d), \beta(d))$ solves

$$
\min_{(\alpha,\beta)\in\sigma^d} E_t\left(u_q(t, \alpha(\beta)) : d\right). \tag{13}
$$

Finally, we argue that neither the principals nor the agents have incentives to deviate from the equilibrium strategies. First, if Principal j offers contract $c_j^{i,\gamma_1,\ldots,\gamma_n}$ for all (j, i), then no player can profitably deviate in the subgame generated by the ordinary contract profile d^* because $(\alpha^*, \beta^*) \in \sigma^{d^*}$. In fact, we have defined the strategies of the players, $\alpha(d)$ and $\beta(d)$, in any relevant subgame generated by an ordinary contract profile, d, such that $(\alpha, \beta) \in \sigma^d$. Therefore, we only have to show that no principal can profitably deviate at the first stage of the game. Recall that if Principal q offers the contract c_q^i instead of $c_q^{i,\gamma_1,\dots,\gamma_n}$ for some i, then his payoff is (13). Hence, the maximum payoff he can achieve by deviating from his equilibrium contract is

$$\max_{d_q\in D_q}\min_{(\alpha,\beta)\in\sigma^{(d_q,d_{-q,q})}} E_t\left(u_q\left(t,\alpha\left(\beta\right)\right):\left(d_q,d_{-q,q}\right)\right).$$

By (10), the previous expression can be rewritten as

$$\min_{d_{-q}\in D_{-q}}\max_{d_q\in D_q}\min_{(\alpha,\beta)\in\sigma^{(d_q,d_{-q})}} E_t\left(u_q\left(t,\alpha\left(\beta\right)\right):\left(d_q,d_{-q}\right)\right)=\underline{u_q}.$$

This implies that Principal q can achieve at most $\underline{u_q}$ by deviating at the first-stage. Therefore, by (ii) of the hypothesis of the theorem, no deviation is profitable. ■

Next, we turn our attention to the more difficult "only if" part of the proof. Let G_d denote the subgame generated by the ordinary contract profile $d\in D$.

Definition 4 *The subgames G_d and $G_{d'}$ $(d,d'\in D)$ are said to be equivalent, $G_d\sim G_{d'}$, if the set of equilibrium outcomes are the same in the two subgames.*[11]

The next proposition states that for all $c_{-j}\in C_{-j}$ there exists a $d_{-j}\in D_{-j}$ such that for all $d_j\in D_j$, Principal j can write a contract so that the subgame generated by the contract profile is equivalent to $G_{(d_j,d_{-j})}$. That is, no matter what the equilibrium contracts are, there always exists an ordinary contract profile d_{-j}, such that Principal j can induce a subgame $G_{(d_j,d_{-j})}$ for all d_j by an appropriate deviation. This implies that it is without loss of generality to assume that the contractual punishment for any deviation by Principal j is simply d_{-j}. That is, the punishment does not depend on the deviation itself, only on the identity of the deviator.

To state this result formally, for all $c=(c_1,\dots,c_n)\in C$ let $d(c)\in D$ denote the ordinary contract profile generated by c. That is, $d_j(c)=c_j([c_1],\dots,[c_n])$ for all $j\in\{1,\dots,n\}$.

Proposition 1 *Let $c=(c_1^*,\dots,c_n^*)\in C$. Then, for all j there exists a $d_{-j}\in D_{-j}$, such that for all $d_j\in D_j$ there exists a $c_j\in C_j$ such that $G_{(d_j,d_{-j})}\sim G_{d\left(c_j,c_{-j}^*\right)}$.*

Proof. See the Appendix. ■

This proposition is key to the "only if" part of the theorem. Since the proof of the proposition is lengthy and technical, it is relegated to the Appendix. Here, we sketch the proof for the case where there are two principals and no agents. Since there are no agents, and therefore the restrictions on the action spaces cannot depend on the messages, a contract of a principal is just a definable mapping from the codes of the contracts to the subsets of the codes of the action space of the

[11] Whether or not two subgames are equivalent depends on the particular equilibrium concept. However, it will become clear from the way this definition is used that our results do not depend on the refinement concept.

principal. Similarly, an ordinary contract is a subset of the codes of the action space of a principal. For all $d_1 \in D_1$, define

$$S(d_1) = \{d_2 : \exists c_1 \ c_1([c_1], [c_2^*]) = d_1, \ c_2^*([c_1], [c_2^*]) = d_2\}.$$

That is, $S(d_1)$ is the set of those values of d_2 for which Principal 1 is able to offer a contract such that the generated subgame is $G_{(d_1,d_2)}$. The statement of the proposition is equivalent to $\cap_{d_1 \in D_1} S(d_1) \neq \{\emptyset\}$. Suppose by contradiction that $\cap_{d_1 \in D_1} S(d_1) = \{\emptyset\}$. This implies that for all $d_2 \in D_2$ there exists a d_1 such that $d_2 \notin S(d_1)$. Therefore, one can construct a function, f $:D_2 \to D_1$, such that $d_2 \notin S(f(d_2))$. Since D_1 and D_2 are finite sets, the function f is definable.[12] Consider now the following contract in one free variable:

$$c_1^x([c_2]) = f\left(c_2^*\left(\left[< x >^{(x)}\right]\right)\right).$$

Let γ denote the Gödel code of this contract. Then $c_1^\gamma([c_2]) = f(c_2^*([c_1^\gamma]))$. Notice that, by the definition of the function f, $c_2^*([c_1^\gamma]) \notin S(f(c_2^*([c_1^\gamma])))$. Substituting the previous equality into $f(c_2^*([c_1^\gamma]))$ we get

$$c_2^*([c_1^\gamma]) \notin S(c_1^\gamma([c])). \tag{14}$$

On the other hand, by the definition of S, $c_2([c_1]) \in S(c_1([c_2]))$ for all c_1,c_2 . Therefore,

$$c_2^*([c_1^\gamma]) \in S(c_1^\gamma([c_2^*])). \tag{15}$$

Of course, (14) and (15) cannot be true simultaneously, and hence, $\cap_{d_1 \in D_1} S(d_1) \neq \{\emptyset\}$.

This result is stated in Lemma 3 of Peters and Szentes (2012) for environments where there are no agents. The difficulty of generalizing this argument for the case when there are agents is that the ordinary contract space of Principal j, D_j, is not finite. Therefore, the function f is not necessarily definable. The proof in the Appendix takes advantage of the fact that although these spaces are infinite, the range of any ordinary contract is finite.

Proof of the "only if" part of Theorem 1. We have already established in the text before the statement of the theorem that g is subgame-implementable. We only have to show that the payoff of Principal j in every equilibrium is at least $\underline{u_j}$ for all $j \in \{1, ..., n\}$. Suppose that $(c_1^*, ..., c_n^*) \in C$ is an equilibrium contract profile. According to Proposition 1 there exists a $d'_{-j} \in D_{-j}$ such that Principal j can generate a subgame which is equivalent to $G_{\left(d_j, d'_{-j}\right)}$ for all $d_j \in D_j$. Let β^* and α^* denote the second-stage equilibrium strategies of the agents and the third-stage equilibrium strategies of the principals, respectively. Then Principal j's equilibrium payoff is weakly larger than

$$\begin{aligned}
\max_{d_j \in D_j} E_t\left(u_j(t, \alpha^*(\beta^*)) \mid \left(d_j, d'_{-j}\right)\right) &\geq \max_{d_j \in D_j} \min_{(\alpha,\beta) \in \sigma\left(d_j, d'_{-j}\right)} E_t\left(u_j(t, \alpha(\beta)) \mid \left(d_j, d'_{-j}\right)\right) \\
&\geq \min_{d_{-j} \in D_{-i}} \max_{d_j \in D_j} \min_{(\alpha,\beta) \in \sigma\left(d_j, d_{-j}\right)} E_t\left(u_j(t, \alpha(\beta)) \mid (d_j, d_{-j})\right) = \underline{u_j}.
\end{aligned}$$

[12] The sets D_1 and D_2 are finite because there are no agents. Therefore an ordinary contract is a restriction on the action space. There are only finitely many such restrictions because the action space of each principal is finite.

■

Note that a principal can offer contracts which do not vary with the contracts of the other principals even in the contractible contracting game. These contracts are effectively ordinary contracts. In fact, if each Principal $q\,(\neq j)$ offers ordinary contracts then the best response of Principal j is also to offer ordinary contracts. Hence, the ability to contract on contracts expands the set of implementable allocations.

Remark 2 *The set of implementable allocations in the contractible contracting game is larger than in the ordinary contracting game.*

Proof. See the Appendix. ■

Section 5 describes an environment where the set of allocations implementable by the contractible contracting game is strictly larger than the set of allocations implementable by ordinary contracts (see Proposition 3).

4.3 The Minmax Values in Special Cases

Whether an allocation is subgame-implementable only depends on the preferences of the principals and the agents. Hence part (i) of the statement of Theorem 1 is a property of an allocation which depends only on the physical environment. However, the minmax values of the principals are defined in terms of equilibria in subgames of the ordinary contracting game. It is desirable to characterize even these minmax values in terms of the physical environment. Next, we show that one additional assumption leads to such a characterization.

Assumption 1. For all j there exist $\overline{a}_j \in A_j$, $a_{-j,j} \in A_{-j}$, and $\mathcal{U}_j : T \to \mathbb{R}$ such that
(i) $u_j\left(t, \overline{a}_j, a_{-j}\right) \geq \mathcal{U}_j\left(t\right)$ for all $a_{-j} \in A_{-j}$, and
(ii) $u_j\left(t, a_j, a_{-j,j}\right) \leq \mathcal{U}_j\left(t\right)$ for all $a_j \in A_j$.

This assumption is satisfied in many important economic applications. The action $\overline{a}_j$ can often be thought of as a default action of Principal j which allows him not to participate in the interaction with the agents. If the principals are sellers and the agents are buyers then $\overline{a}_j$ means that Principal j does not sell his products. If the principals are employers and the agents are workers then this action corresponds to the choice of not employing any worker. The action profile $a_{-j,j}$ can be interpreted as an action profile of the principals (other than Principal j) which excludes Principal j from participation. In the buyer-seller example, this can be accomplished by setting prices so low that Principal j cannot make a positive profit by selling his products. Similarly, in the employer-worker example, the principals can set wages higher than the productivity of the workers.

Theorem 2 *Suppose that Assumption 1 is satisfied for the mappings $\mathcal{U}_j$ for all $j \in \{1, ..., n\}$. Then the allocation $g : T \to A$ is implementable as an equilibrium in the contractible contracting game if and only if (i) $g \in \mathcal{A}$, and (ii) $E_t u_j\left(t, g\left(t\right)\right) \geq E_t \mathcal{U}_j\left(t\right)$ for all $j \in \{1, ..., n\}$.*

Proof. See the Appendix. ■

Next, we show that if information is complete, our Theorem 1 also leads to a characterization of the equilibria in terms of the physical environment without any reference to ordinary contracting games.

Theorem 3 *Suppose that $|T^i| = 1$ for all i. Then the allocation $(a_1^*, ..., a_n^*) = a^* \in A$ is implementable as a subgame perfect Nash equilibrium if and only if*

$$u_j(a^*) \geq \min_{a_{-i}} \max_{a_i} u_j(a_i, a_{-i}) = u_j^*. \tag{16}$$

A notable feature of this corollary is that whether or not an allocation is implementable does not depend on the number or preferences of the agents.

Proof. See the Appendix. ■

5 Application to Exclusive Contracting: Welfare and Policy Implications

Arguably, in many real-life applications of common-agency models, an agent signs a contracts with only one of the principals. Examples include employer-worker relationships in which the worker can accept only one job offer, and seller-buyer relationships in which the buyer is interested in purchasing a good or a service from only one seller. In this section we consider such environments, which we refer to as *exclusive-contracting environments* (formally defined below).

We begin by showing that, in exclusive-contracting environments, strong incentive compatibility coincides with the standard notion of incentive compatibility. We prove that any allocation which is incentive compatible and individually rational can be implemented in the contractible contracting game. We next direct our attention to settings in which the principals have identical action spaces and identical payoff functions. In the buyer-seller example, this implies that the sellers' products are perfect substitutes from the agent's point of view, and that all sellers have the same production costs. We refer to this quite naturally as a *perfectly competitive environment.* By restricting attention to perfectly competitive, exclusive-contracting environments, we are able to define and compare monopolist and competitive allocations. Indeed, as both the action spaces and payoffs are identical across principals, simply varying the number of principals becomes a meaningful exercise. We show that if there is more than one principal and contracts are not contractible, then every equilibrium outcome is Pareto efficient. However, if contracts are allowed to be contractible, any efficiency gain generated by the competition may disappear due to collusion between principals. In fact, through collusion, principals are even able to implement the monopolist outcome. Therefore, as a policy implication, our results suggest the prohibition of contracting on contracts in this type of environment.

Exclusive Contracting.— There are many principals ($n > 1$) and a single agent ($k = 1$). The finite type space of the agent is T. If Principal j ($= 1, ..., n$) and the agent enter into a contractual relationship, the possible *agreements* between them can be described by the set $X \times P$, where X is a finite set of contractible decisions and $P \subset \mathbb{R}$ is a finite set of transfers.[13] Contracting is exclusive: the agent can enter into a contractual agreement with only one principal. We emphasize that this exclusivity assumption still allows for the terms of a principal's contract to depend on the contracts of the other principals. If the agent with type t signs a contract with Principal j and the contract specifies $(x_j, p_j) \in X \times P$, then the payoffs to the agent and Principal j are

$$V(t, x_j) - p_j \text{ and } U_j(t, x_j) + p_j,$$

respectively. The agent's outside option is normalized to be zero. That is, if the agent chooses not to sign a contract he receives a payoff of zero. Similarly, if Principal j's contract is not signed by the agent, she receives a payoff of zero.

In order to simplify the discussion of the principals' and the agent's outside options, we assume that $0 \in X \cap P$ for all $j = 1, ..., n$ and $V(t, 0) = U_j(t, 0) = 0$ for all $t \in T$. As a consequence, it is without loss of generality to restrict attention to allocations in which the agent chooses to participate, and signs a contract with a principal. Formally, an allocation in the exclusive-contracting environment is a triple, $(\mathbf{x}, \mathbf{p}, \mathbf{n})$, $\mathbf{x} : T \to X$, $\mathbf{p} : T \to P$, $\mathbf{n} : T \to \{1, ..., n\}$ such that $V(t, \mathbf{x}(t)) - \mathbf{p}(t) \geq 0$ for all $t \in T$. The interpretation of this triple is as follows: the agent with type t contracts with Principal $\mathbf{n}(t)$ and the payoffs are determined by the decision-transfer pair $(\mathbf{x}(t), \mathbf{p}(t))$.

This model has various possible interpretations. For example, one might think of the agent as a buyer and the principals as sellers. The agent's type is his valuation, x_j is the quality (or quantity) of the principal's product, and p_j is the price. If instead, the agent is a worker and the principals are employers, then the pair (x_j, p_j) is a labor contract where p_j is the wage and x_j specifies other characteristics of the job such as working hours, the number of vacation days, or health insurance benefits. A worker's type might correspond to his productivity, or his taste for various characteristics of the job. As a third example, suppose the agent is an economic consultant and the principals are competing firms in a certain industry. The type of the consultant is his industry experience, and the contractible decision is the deadline by which the consultant must complete the project. Of course, many other sensible interpretations exist.

Transformation.— The exclusive-contracting environment differs from the original model in Section 3 in two ways. First, in our original model, the agent entered into a contractual agreement with each principal, while in the exclusive-contracting environment the agent must choose a single principal. Second, in the original model, payoffs are defined as a function of the action profiles of the principals; in the exclusive-contracting environment, payoffs depend only on the action of

[13] Note that it is without loss of generality that the set X is the same across principals, because a principal's inability to make a certain decision can be captured by assigning to her a large negative payoff.

the principal whose contract is accepted by the agent. Nevertheless, we are able to demonstrate that an exclusive-contracting environment can be transformed into a strategically equivalent model which is a special case of our original model. The main idea is to encode the agent's choice of principal into the definition of the payoffs. To this end, we define the payoffs of the principals and the agent as a function of the agreement profiles. In order to be consistent with the notation in Section 3, let A_j denote $X \times P$ for all $j = 1, ..., n$. If Principal j takes action $a_j = (x_j, p_j)$, then the agent's payoff function, $v : T \times A \to \mathbb{R}$, is defined by

$$v(t, a) = \max\left\{0, \max_{j \in \{1,...,n\}} V(t, x_j) - p_j\right\}. \tag{17}$$

That is, the agent's payoff is determined by the agreement offered by whichever principal can make him best off, unless he opts out and receives his reservation value, which is zero. In order to resolve ties, we assume that if $V(t, x_j) - p_j = V(t, x_q) - p_q$ and $j < q$ then the agent strictly prefers (x_j, p_j) to (x_q, p_q). The payoff to Principal j is defined as follows:

$$u_j(t, a) = \begin{cases} u(x_j) + p_j & \text{if } j = \min \arg\max_{q \in \{1,...,n\}} (V(t, x_q) - p_q) \text{ and } V(t, x_j) - p_j \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

A principal is only able to obtain a positive payoff if her proposed agreement maximizes the agent's payoff. In addition, her name must be the first one listed among those who manage to maximize the agent's payoff. Let us point out that the payoff functions $\{u_j\}_{j=1}^n$ and v have the same domain as those in Section 3. We also note that the agent's choice of the principal is encoded in the definitions of v and $\{u_j\}_1^n$. Therefore, the exclusive-contracting environment can indeed be treated as a special case of our general model.

Next, we characterize the allocations which are implementable by contractible contracting games.

Definition 5 *The allocation* $(\mathbf{x}, \mathbf{p}, \mathbf{n})$, $\mathbf{x} : T \to X$, *and* $\mathbf{p} : T \to P$ *is incentive compatible and individually rational if*

$$\begin{aligned} V(t, \mathbf{x}(t)) - \mathbf{p}(t) &\geq V(t, \mathbf{x}(t')) - \mathbf{p}(t') \text{ for all } t, t' \in T, & (18) \\ E_t(U_j(t, \mathbf{x}(t)) + \mathbf{p}(t) \mid \mathbf{n}(t) = j) &\geq 0 \text{ for all } j = 1, ..., n. & (19) \end{aligned}$$

The inequalities in (18) represent the agent's incentive compatibility constraints, and the inequalities in (19) are the principals' participation constraints. Recall that the agent's participation constraint is included in the definition of an allocation, that is, $V(t, \mathbf{x}(t)) - \mathbf{p}(t) \geq 0$ for all t.

In what follows, we assume that the range of transfers is sufficiently large, in particular, $V(t, x) - U_j(t, x) \leq -\min P$ for all j, x and t.

Proposition 2 *In the contractible contracting game, an allocation can be implemented if and only if it is incentive compatible and individually rational.*

It is worth comparing the statement of this proposition with those of Theorems 1 and 2. One notable difference is that, in an exclusive contracting environment, strong incentive compatibility of an allocation is no longer a requirement for impelementability. Indeed, in the proof we argue that the set of strong incentive compatible allocations coincides with that of incentive compatible allocations.

Proof. First, we argue that Assumption 1 holds for the exclusive-contracting environment and hence, Theorem 2 is applicable. For each $j = 1, ..., n$ and $t \in T$, define $\mathcal{U}_j(t)$ to be zero. In addition, let $\overline{a}_j = (0, 0)$ for all j and $a_{-j,j} = (0, \min P)$ for all j. The triple $\left(\{\mathcal{U}_j\}_1^n, \{\overline{a}_j\}_1^n, \{a_{-j,j}\}_1^n\right)$ clearly satisfies parts (i) and (ii) of Assumption 1. Therefore, Theorem 2 and Remark 1 imply that an allocation is implementable if and only if it is strongly incentive compatible and Principal j receives a payoff of at least $E_t\mathcal{U}_j(t)$ for all $j \in \{1, ..., n\}$. Since $\mathcal{U}_j(t) \equiv 0$, these latter conditions coincide with the inequalities in (19). Hence, in order to complete the proof, it remains to show that the set of strongly incentive compatible allocations, $\mathcal{A}$, is the same as the set of incentive compatible allocations.

Recall that strong incentive compatibility implies incentive compatibility. So, we only need to show the reverse: that in an exclusive-contracting environment, each incentive compatible allocation is also strongly incentive compatible. Note that the allocation$(\mathbf{x}, \mathbf{p}, \mathbf{n})$ in the exclusive contracting environment corresponds to the allocation $g = (g_1, ..., g_n) : T \to A$, such that

$$g_j(t) = \begin{cases} (\mathbf{x}(t), \mathbf{p}(t)) & \text{if } \mathbf{n}(t) = j \\ (0, 0) & \text{otherwise.} \end{cases} \tag{20}$$

According to Definition 3, we need to show that (18) implies that for all $t \in T$, and $(t_1, ...t_n) \in (T)^n$:

$$v(t, (g(t))) \geq v(t, (g_1(t_1), ..., g_n(t_n))).$$

By (17) and (20), this inequality can be rewritten as

$$V(t, \mathbf{x}(t)) - \mathbf{p}(t) \geq \max_{t'} [V(t, \mathbf{x}(t')) - \mathbf{p}(t')],$$

which follows from (18). ■

Perfectly Competitive Environments.— Next, let us consider environments in which the principals compete, that is, each principal receives the same payoff when making a particular decision. To be more specific, we assume that $U_j(t, x) = U(x)$ for all $j = 1, ..., n$, $x \in X$ and $t \in T$. The assumption that the principal's payoff is not directly affected by the agent's type is made in order to guarantee the existence of an equilibrium in a competitive screening model, and is satisfied in most of the applications mentioned above.[14] As is shown below, this requirement enables us to provide a simple characterization of competitive equilibria.

[14] Perhaps the most prominent example in which an equilibrium does not exist is the insurance model of Rothschild and Stiglitz (1976). For a detailed discussion of equilibrium existence in competitive screeining models, see Chapter 13 of Mas-Colell et al. (1995).

Consider the problem of a single (monopolist) principal. The monopolist maximizes $E_t\left[U\left(t, \mathbf{x}\left(t\right)\right) + \mathbf{p}\left(t\right)\right]$ subject to (18) and the participation constraint of the agent. We denote the solution to this problem by $\left(\mathbf{x}_m\left(t\right), \mathbf{p}_m\left(t\right)\right)$ and refer to it as the *monopoly outcome*. Since $U\left(0\right) = 0$, the participation constraint of the principal in (19) is automatically satisfied at the solution to this maximization problem. The Pareto Optimal outcome solves the following problem:

$$\max_{\mathbf{x}} U\left(\mathbf{x}\left(t\right)\right) + V\left(t, \mathbf{x}\left(t\right)\right) \text{ for all } t \in T.$$

Let $\mathbf{x}_c$ denote the solution.

We assume that the set of transfers, P, is rich enough to implement the Pareto Optimal outcome. That is, for all t there is a $p_t \in P$ such that $V\left(t, \mathbf{x}\left(t\right)\right) \geq p_t \geq -U\left(\mathbf{x}\left(t\right)\right)$. We also assume that the range for transfers is sufficiently large, that is, $U\left(x\right) + \min P < 0 < U\left(x\right) + \max P$ for all $x \in X$. These assumptions are only needed because we are considering a discrete model, and are satisfied in the standard continuous version of our environment. Note that $\left(\mathbf{x}_m, \mathbf{p}_m\right)$ as well as $\mathbf{x}_c$ are generically unique.

Proposition 3 *Suppose that an exclusive-contracting environment is perfectly competitive. If* $n > 1$*, each equilibrium implements* $\mathbf{x}_c$ *in the ordinary contracting game. In the contractible contracting game, there exists an equilibrium which implements* $\mathbf{x}_m$*.*

According to this proposition, in perfectly competitive, exclusive environments, the contractibility of contracts can offset any efficiency gain generated by competition among the principals. The policy implication is that contracting on contracts should be prohibited in these environments.

An implication of this proposition is that the set of implementable allocations in the contractible contracting game is strictly larger than in the ordinary contracting game. Note that this proposition can be easily generalized for the case of $k\left(> 1\right)$ identical agents as long as the payoffs of the principals are additive in the agents. Therefore, we conclude that allowing the contractibility of contracts might strictly enlarge the set of equilibrium outcomes, irrespective of the number of agents.

Proof. Suppose that contracts are not contractible and, by contradiction, that there exists an equilibrium and $t \in T$ such that if the agent's type is t, $\mathbf{x}_c\left(t\right)$ is not implemented. Since the agent prefers to contract with the principal with the *smallest* name, all but Principal 1 receive a payoff of zero, even conditional on the type of the agent. (This is because if the agent accepts Principal j's $\left(j \neq 1\right)$ contract then Principal 1 could offer the same contract and achieve the same payoff as Principal j.) We show that Principal 2 can increase his payoff by deviation at the contracting stage. Consider the following contract: $d\left(l\right) = \left(\mathbf{x}_c\left(t\right), \mathbf{p}_c\left(t\right)\right)$ for all l. By part (i) of Assumption 2, the agent with type t, and perhaps with other types too, will interact with Principal 2, and hence, Principal 2 can achieve a positive payoff. In addition, this payoff is strictly positive generically.

Suppose now that contracts are contractible. In order to verify that $\left(\mathbf{x}_m, \mathbf{p}_m\right)$ is an equilibrium outcome, define $\mathbf{n}\left(t\right) \equiv 1$. We argue that the allocation $\left(\mathbf{x}_m, \mathbf{p}_m, \mathbf{n}\right)$ is incentive compatible and individually rational. First, note that all principals except Principal 1 receive a payoff of zero. Second, Principal 1 receives the monopoly profit which is weakly positive because $U(0) = 0$. Finally, $(\mathbf{x}_m, \mathbf{p}_m, \mathbf{n})$ satisfy both the agent's incentive and participation constraints because $(\mathbf{x}_m, \mathbf{p}_m)$ is the solution to the monopolist's maximization problem subject to the exact same constraints. Therefore, Proposition 2 implies that $(\mathbf{x}_m, \mathbf{p}_m, \mathbf{n})$ is an equilibrium outcome in the contractible contracting game. ■

Next, we show that if the environment is not perfectly competitive, the contractibility of the contracts can in fact lead to an increase in welfare.

Example 2. Suppose that $n = 2$, $k = 1$, $X = \{0, a, b\}$ and $T = \{(A, H), (A, L), (B, H), (B, L)\}$. We shall interpret this example in the following way: there are two firms (principals) and each of them is specialized in producing a distinct product: Firm 1 produces good a and Firm 2 produces good b. The buyer's (agent's) type is two-dimensional, and he is interested in buying at most one unit of a good. The first dimension corresponds to the product the buyer is interested in buying while the second dimension is his valuation, which is either high (H) or low (L). The following table describes the utility profile (U_1, U_2, V) for each $(x, t) \in X \times T$:

	(A,H)	(A,L)	(B,H)	(B,L)
a	$0, -\infty, H$	$0, -\infty, L$	$0, -\infty, 0$	$0, -\infty, 0$
b	$-\infty, 0, 0$	$-\infty, 0, 0$	$-\infty, 0, H$	$-\infty, 0, L$

,

The first and second numbers in each cell denote the payoffs to Principal 1 and Principal 2 respectively, and the third number is the payoff to the agent. Note that if the agent's type is A (B), his willingness-to-pay for the other good b (a) is zero, regardless of whether his valuation is high or low. In this example, U_1 and U_2 are the principals' costs of production. Note that Principal 1's (2's) inability to produce good b (a) is captured by assigning to her an infinitely large production cost. Finally, we assume that $H > 2L > 0$.

Since the buyer's willingness-to-pay is positive for one of the goods, the Pareto Optimal outcome is defined by

$$(\mathbf{x}_c(t), \mathbf{n}_c(t)) = \begin{cases} (a, 1) & \text{if } t \in \{(a, L), (a, H)\}, \\ (b, 2) & \text{if } t \in \{(b, L), (b, H)\}. \end{cases}$$

In what follows, we show that the Pareto Optimal outcome can only be implemented in the contractible contracting game. To this end, let $\mathbf{p}_c(t)$ equal zero for all $t \in T$, that is, the firms sell their products at a price of zero. The allocation $(\mathbf{x}_c, \mathbf{p}_c, \mathbf{n}_c)$ obviously satisfies both (18) and (19) hence, it is incentive compatible and individually rational. Therefore, Proposition 2 implies that $(\mathbf{x}_c, \mathbf{p}_c, \mathbf{n}_c)$ is indeed an equilibrium outcome of the contractible contracting game.

Suppose now that contracts are not contractible. Note that Principal 1 (2) can only make a positive profit if the first dimension of the buyer's type is A (B). Therefore, the profit-maximizing contract of Principal 1 (2) is the monopoly offer conditional on the first dimension of the agent's type being A (B). Since $H > 2L$, Principal 1 (2) offers the good a (b) at price H, and the agent buys the good only if his valuation is H. Hence, the unique equilibrium outcome is

$$(\mathbf{x}(t), \mathbf{p}(t), \mathbf{n}(t)) = \begin{cases} (a, H, 1) & \text{if } t = (a, H), \\ (b, H, 2) & \text{if } t = (b, H), \\ (0, 0, 0) & \text{otherwise.} \end{cases}$$

This allocation is clearly inefficient: an agent with valuation L leaves empty-handed despite the fact that his willingness to pay exceeds the cost of production.

In the previous example, two firms produced different products and the demand for their products came from different consumers. Therefore, in the ordinary contracting game, each firm acts as a monopolist and generates a deadweight loss. If contracts are contractible, each firm can promise a large payment to the consumer unless the other firm adopts competitive pricing. Of course, the other firm's best response is to set the competitive price. In this way, efficiency can be restored in the contractible contracting game. Note, however, that the efficient equilibrium requires that two firms engaging in very different activities to contract on each others' contracts. This would mean, for example, that a bank's wage offer to a financial analyst would depend explicitly on the wage contract offered to a general manager of a sport club. Such behavior seems both unrealistic and absurd, at least from an applied point of view. Let us emphasize that Example 2 merely illustrates the importance of the competitiveness assumption in the statement of Proposition 3.

6 Attributes of the Equilibrium Contracts

This section further investigates the attributes of the equilibrium contracts in both the ordinary and the contractible contracting games.

6.1 Examples for Ordinary Contracting Games

For simplicity, we identify the message of an agent with its Gödel code in all the examples below. That is, instead of saying that an agent sends a message whose Gödel code is q, we say that the agent sends the message q. (This does not cause confusion because the encoding is a bijection.)

Next, we show, by examples, that one cannot assume that the equilibrium contracts specify a single action for a principal as a function of the agent's message. The contract d_j^i is said to be *complete* if $|d_j^i(q)| = 1$ for all q, that is, d_j^i is a function from $\mathbb{N}$ to A_j^i. Restricting the contracts to be complete is with the loss of generality for two reasons. Example 3 shows that there are allocations which cannot be supported with complete contracts, but can be supported otherwise. Example 4 shows that there are allocations which can only be supported if contracts are required to be complete.

Example 3. Suppose that $n = 2$ and $k = 1$. Assume that the agent's type space is degenerate,

$A_1 = A_2 = \{a, b\}$, and the payoffs are defined by the following matrix:

	a	b
a	2, 2, 0	0, 3, 3
b	1, 0, 0	1, 0, 1

,

where the first and second numbers in each cell describes the payoffs to Principal 1 and Principal 2, and the third number is the payoff to the agent.

Notice that the agent's payoff is zero whenever Principal 2 takes action a and positive otherwise. Therefore, whenever he can send a message which triggers action b by Principal 2, he will do so. In addition, given that Principal 2 takes action b, the agent prefers Principal 1 to take action a over action b. Consider the allocation (a, a). Principal 2 would like to deviate and take action b. Such a deviation can be punished by Principal 1 by taking action b. We show that the outcome (a, a) can be implemented as an equilibrium but cannot be implemented with complete contracts.

Define the equilibrium contracts of the principals as follows: $d_1(q) = A_1$ for all q, and $d_2(q) = \{a\}$ for all q. Since these contracts are constants in the messages of the agents, the strategy of the agent is irrelevant. Principal 1's strategy is the following. If he observes that Principal 2 offered a contract which allows taking action b for some reports of the agent, he takes action b, otherwise he takes action a. Obviously, none of the principals can increase his payoff by offering a different contract.

Next, we argue that (a, a) cannot be supported by complete contracts. Suppose that (d_1, d_2) supports (a, a) and d_1 is complete. Then, there exists an $q \in \mathbb{N}$ such that $d_1(q) = \{a\}$. Then Principal 2 can profitably deviate by offering a contract which specifies action b independently of the agent's report. This is because the agent reports a $q \in \mathbb{N}$ to Principal 1 such that $d_1(q) = \{a\}$ and the outcome will be (a, b). This outcome maximizes the agent's payoff and provides Principal 2 with a payoff higher than the outcome (a, a) would.

Example 4. Suppose $n = 2$ and $k = 1$. Assume that the agent's type space is degenerate, $A_1 = A_2 = \{H, T\}$, and the payoffs are defined by the following matrix:

	H	T
H	1, -1, -1	-1, 1, 1
T	-1, 1, 1	1, -1, -1

,

where the first and second numbers in each cell describe the payoffs of Principal 1 and Principal 2, and the third number is the payoff to the agent. In this example, the two principals are playing the Matching Pennies Game, and the agent's payoff is identical to that of Principal 2.

We first show that if each principal is restricted to offer a complete contract then the payoff profile $(-1, 1, 1)$ can be supported as an equilibrium payoff profile. To see this, consider the following contract of Principal 2: $d_2(1) = \{H\}$ and $d_2(q) = \{T\}$ if $q \neq 1$. Suppose that the complete contract of Principal 1 is d_1. Notice that $d_1(1)$ is either H or L. If $d_1(1) = \{H\}$ then the agent can send messages 1 and 2 to Principals 1 and 2 respectively, which generates a payoff profile $(-1, 1, 1)$. Similarly, if $d_1(1) = \{L\}$, the agent can send the message 1 to both principals, which again generates a payoff profile of $(-1, 1, 1)$. Therefore, no matter what the complete contract of Principal 1 is, the agent can always induce the payoff profile $(-1, 1, 1)$.

Suppose now that the principals are not restricted to offer complete contracts. Then there does not exist a pure strategy equilibrium in our game, because Principal 1 can always offer a contract d, such that $d(q) = \{H, L\}$. In addition, if we allow mixed strategies, the only equilibrium payoff profile was $(1/2, 1/2, 1/2)$.

Next we show that one cannot assume that the message space of an agent is his type space. To be more specific, the next example shows that the cardinality of the range of the equilibrium contracts must be larger than the cardinality of the type space of the agent in order to implement certain allocations.[15]

Example 5. Suppose that $n = 2$ and $k = 1$ and the type space of the agent is degenerate. The principals are playing the Prisoner's Dilemma. That is, $A_1 = A_2 = \{C, D\}$, and the payoffs are defined by the following matrix

	C	D
C	2, 2, 3	0, 3, 1
D	3, 0, 1	1, 1, 2

Again, the first two numbers are the payoffs to the principals and the third one is the payoff to the agent. Notice that the agent prefers the principals to cooperate to everything else, but prefers them to defect to (C, D) and (D, C). The agent has no private information in this example. Hence, if the action profile (C, C) could be implemented such that the message space of the agent is his type space then (C, C) would be supported as an equilibrium outcome by contracts which do not depend on the report of the agent. We show that this is impossible although (C, C) is implementable.

Suppose that d_1 and d_2 implement (C, C) and $d_i(q_1) = d_i(q_2)$ for all $q_1, q_2 \in \mathbb{N}$ and $i \in \{1, 2\}$. If $d_1(q) = \{C\}$ for all q, Principal 2 can profitably deviate by offering a contract that specifies $\{D\}$. Hence, $d_1(q) = d_2(q) = \{C, D\}$ for all q, which implies that the principals play the Prisoner's Dilemma in the last stage of the game, and therefore, $\{C, C\}$ cannot be implemented.

Now, we show that we can implement $\{C, C\}$ with the help of the agent. Consider the following contract

$$d_i(q) = \begin{cases} C & \text{if } q = 1, \\ D & \text{if } q \neq 1. \end{cases}$$

The strategy of the agent is defined such that he triggers (C, C) whenever he can. In particular, on the equilibrium path, the agent reports 1 to each principal. The agent has no incentive to deviate because his payoff is maximized. If one of the principals deviates and offers a contract such that the agent cannot induce the action C, the agent reports 2 to the other principal, and the outcome would be $\{D, D\}$.

[15] Similar results appear in Peck (1997), Martimort and Stole (2002), and Attar, Mariotti and Salanie (2011).

6.2 Incompleteness of Contracts in the Contractible Contracting Games

Recall that Example 3 showed that contracts cannot be assumed to be complete in the ordinary contracting game. The following example shows that this is true even if contracts are contractible.

Example 6. Suppose that $n = 2$ and $k = 1$. The action space of Principal 1 is $\{x, y\}$, and the action space of Principal 2 is $\{a_1, a_2, s\}$. The type space of the agent is $\{1, 2\}$, and each type is equally likely. The payoff of Principal 1 is constant zero. The following tables represent the payoffs to Principal 2 and to the agent, respectively:

$t = 1$	a_1	a_2	s
x	$1, 1.1$	$-3, 0.1$	$0, -1$
y	$0, 1$	$-3, 0$	$1, -1$

,

$t = 2$	a_1	a_2	s
x	$-3, 0$	$1, 1.1$	$0, -1$
y	$-3, 0.1$	$0, 1$	$1, -1$

First, we show that the constant allocation (x, s) can be implemented with incomplete contracts. Then we show that the same allocation cannot be implemented with complete contracts.

In order to show that the outcome (x, s) is an equilibrium outcome, we have to verify the two conditions of Theorem 1. The constant action profile (x, s) belongs to $\mathcal{A}$ because it is the unique outcome in the subgame generated by (d_1, d_2), where $d_1(n) = \{x\}$ and $d_2(n) = \{s\}$ for all n. It remains to show that the minmax value of Principal 2, defined by (9), is weakly smaller than zero. In order to do so, consider the following contract of Principal 1: $d_1(n) = \{x, y\}$ for all $n \in \mathbb{N}$. We show that for all $d_2 \in D_2$, there is an equilibrium in the subgame $G_{(d_1,d_2)}$ such that the payoff of Principal 2 is at most zero. We have to analyze four different cases depending on the range of d_2.

Case 1. There exist n_1 and n_2 such that $d_2(n_1) = \{a_1\}$ and $d_2(n_2) = \{a_2\}$. Define the agent's strategy as follows: $m_1^2(1) = n_1$, $m_1^2(2) = n_2$, and $m_1^1 \equiv 1$. Principal 1's strategy is to take action y. These strategies constitute an equilibrium in the subgame $G_{(d_1,d_2)}$ and result in a payoff of 0 to Principal 2.

Case 2. There exists an n_1 such that $d_2(n_1) = \{a_1\}$, but there does not exist n_2 such that $d_2(n_2) = \{a_2\}$. Define the agent's strategy such that $m_1^2 \equiv n_1$ and $m_1^1 \equiv 1$. Principal 2's strategy is defined as follows. If $s \in d_2(m_1^2)$, then he takes action s, otherwise he takes action a_1. Principal 1 always takes action x. These strategies constitute an equilibrium in the subgame $G_{(d_1,d_2)}$ and result in an expected payoff of minus two to Principal 2.

Case 3. There exists an n_2 such that $d_2(n_2) = \{a_2\}$, but there does not exist any n_1 such that $d_2(n_1) = \{a_1\}$. Define the agent's strategy such that $m_1^2 \equiv n_2$ and $m_1^1 \equiv 1$. Principal 2's strategy is defined as follows. If $s \in d_2(m_1^2)$, then he takes s, otherwise he takes action a_2. Principal 1 always takes action x. Again, these strategies constitute an equilibrium in the subgame $G_{(d_1,d_2)}$ and result in an expected payoff of minus two to Principal 2.

Case 4. Suppose that there does not exist any n_i such that $d_2(n_i) = \{a_i\}$ for $i = 1, 2$. Define the agent's strategy such that $m_1^2 \equiv n$, if there exists an n such that $d_2(n) = \{a_1, a_2\}$. Otherwise, $m_1^2 \equiv 1$. In addition, $m_1^1 \equiv 1$. Principal 2's strategy is defined as follows. If $s \in d_2(m_1^2)$, then he takes s, otherwise he takes action a_2. Principal 1 always takes action x. These strategies constitute an equilibrium in the subgame $G_{(d_1,d_2)}$ and the payoff of Principal 2 is at most zero.

Now we show that (x,s) cannot be implemented with complete contracts. Suppose by contradiction that there is an equilibrium implementing (x,s), and Principal 1's contract is complete. Let d_1^2 denote the ordinary contract used by Principal 1 to punish deviations of Principal 2. (The existence of such an ordinary contract is guaranteed by Proposition 1.) The contract d_1^2 is complete because Principal 1's equilibrium contract is complete. We have to consider two cases. Case 1: There exists an n such that $d_1^2(n) = \{x\}$. Then, consider the following ordinary contract of Principal 2 : $d_2(1) = \{a_1\}$ and $d_2(n) = \{a_2\}$ if $n \neq 1$. According to Proposition 1, Principal 2 can induce a subgame equivalent to $G_{\left(d_1^2,d_2\right)}$. Then the agent can generate the action profile (x, a_1) if $t = 1$ and the action profile (x, a_2) if $t = 2$. He can do so, for example, by using the reporting strategy defined by $m_1^1 \equiv n$, $m_1^2(1) = 1$ and $m_1^2(2) = 2$. Notice that these action profiles are the unique maximizers of the agent's payoff, hence he will generate this outcome in $G_{\left(d_1^2,d_2\right)}$. But this outcome provides Principal 2 with a payoff of one, which is strictly larger then his payoff from (x,s), and hence, the deviation generating $G_{\left(d_1^2,d_2\right)}$ is profitable. Case 2: Suppose that $d_1^2(n) = \{y\}$ for all n. Then Principal 2 can deviate and generate $G_{\left(d_1^2,d_2\right)}$ where $d_2(n) = s$ for all n. The outcome of this subgame is (y,s), which generates a payoff of one to Principal 2. Hence, Principal 2 can again profitably deviate.

7 Discussion

The goal of this paper was to understand the consequences of the contractibility of contracts in common agency models. Our folk theorem shows that the contractibility of contracts leads to a large set of equilibrium allocations. The interpretation of this theorem is that principals are able to collude and implement various outcomes. At first glance, this result might seem counterintuitive because players do not interact repeatedly. In a repeated environment, a collusion can be sustained because players can punish a deviator in periods followed by a deviation. In our model, the contract of a principal reveals his *intentions*. Therefore, in some sense, the principals observe a deviation even before a payoff-relevant action is taken. Furthermore, since contracts are contractible, the principals can commit to punish deviations, even if they would find it suboptimal ex post.

One might ask why we chose to model the contract space with the set of definable correspondences. Definability played two important roles in our analysis. First, the set of definable functions contains cross-referential functions which were used to construct the equilibrium contracts in our proofs. Second, identifying contracts with definable functions enabled us to pin down the min-max values of the principals. In particular, the statement of Proposition 1 crucially depends on definability.

There are other spaces which contain cross-referential objects. One such a space is, for example, the set of Turing machines which is often used in game theoretic analysis. One can even think about the programs in Tennenholtz (2004) as Turing machines, who used this space in a context similar to others. If we modelled the contracts by Turing machines, then the input would be the descriptions of the machines and the messages. The output of each machine would be a subset of a principal's action space. It is easy to show that the "if part" of our main theorem holds with such a contract space as well. That is, any allocation satisfying the two properties of Theorem 1 can be implemented as an equilibrium.[16] The problem is with the "only if" part of the theorem, and in particular, with the minmax values of the principals. The reason is that principals could submit *universal machines*, which would simulate the machine of a deviator. Once the simulation is completed, these machines could recommend an action profile which is worse for the deviator. This action profile can depend on the result of the simulation, that is, on the actual deviation. This suggests that the principals can push the payoff of a deviator below his minmax value. But of course, a deviator could also submit a universal machine which would simulate the machines of the others, and then best-respond to their outputs. The problem is that, in general, these universal machines will not halt on each other. Indeed, it is not clear how one can properly define the game because of this halting problem.

On might argue that the Gödel coding is an unrealistic feature of our model. Indeed, we do not observe contracts referring to the codes of other contracts. However, as we explained in the introduction, the crucial assumption driving our results is that contracts must written in a language. The Gödel coding itself is definable and hence, it can be embedded into a contract. In fact, players do not even need to agree to use the same codes. They can use any coding unilaterally, and the implications of the contract will be understood by the others, provided that they use the same underlying language. Any language which is rich enough to express arithmetic statements, for example English or Hungarian, does contain cross-referential statements. The set of texts seems to be a natural description of the set of feasible contracts. In order to eliminate the possibility of writing contracts which are cross-referential, one must make restrictions on the texts which are admissible as contracts. It is not clear to us what these natural restrictions should be.

8 Appendix

8.1 Arithmetic Statements

Below, we provide a formal definitions for *arithmetic statement* and arithmetic statements with free variables. We shall define these objects for any first-order logic and explain the specifics of Number Theory

Each formal language has a set of symbols. The symbols of a language are divided into two disjoint sets: the logical-symbols, and the non-logic symbols. The logical-symbols include: (,), $\forall$, $\exists$, $\neg$, $=$, and infinitely many *variable symbols*, $x_0, x_1, ...$ The non-logic symbols include function-symbols and relation-symbols.

[16] In fact, any contract considered in our proof is a recursive function.

Definition 6 *$t = \langle F, R, \tau \rangle$ is a similarity type, where F is a set of function-symbols, R is a set of relation-symbols, and $\tau : F \cup R \to \mathbb{N} \cup \{0\}$ such that $\tau(r) > 0$ if $r \in R$.*

The function τ determines the number of variables of the functions and the relations. If $\tau(f) = 0$, then f is referred to as a *constant-symbol*.

One of the similarity type corresponding to the Peano Arithmetics, denoted by $q = \langle F, R, \tau \rangle$, is: $F = \{0, 1+, *\}$, $R = \{<\}$, $\tau(0) = \tau(1) = 0$, $\tau(+) = \tau(*) = \tau(<) = 2$. Notice that the zero and the one are considered as functions with zero variable, that is, they are constant symbols. (We point out that the similarity type of arithmetics can be defined without the relation "$<$". This relation can be then defined recursively.)

Definition 7 *Let $t = \langle F, R, \tau \rangle$ be a similarity type. Then the set of* expressions *of type t, denoted by $K(t)$, is the smallest set for which:*

(i) $x \in K(t)$ for all variable symbols,

(ii) For all $f \in F$, if $\tau(f) = 0$ then $f \in K(t)$,

(iii) For all $f \in F$, if $\tau(f) = n$, and $k_1, ..., k_n \in K(t)$ then $f(k_1, ..., k_n) \in K(t)$.

Suppose that $t = q$. Then the following string of symbols are expressions in arithmetics: x, 0, 1, $x + 1$, $((x + 1) * (y + 1) + 1)$ etc.

We are ready to define the set of statements corresponding to a similarity type.

Definition 8 *Let $t = \langle F, R, \tau \rangle$ be a similarity type. Then the set of* statements *of type t, denoted by $F(t)$, is the smallest set for which:*

(i) if $r \in R$, $\tau(r) = n$, and $k_1, ..., k_n \in K(t)$ then $r(k_1, ..., k_n) \in F(t)$,

(ii) if $k_1, k_2 \in K(t)$ then $k_1 = k_2 \in F(t)$

(iii) if $\phi, \eta \in F(t)$, then $(\phi) \vee (\eta) \in F(t)$, $\neg(\phi) \in F(t)$, and $\exists x(\phi) \in F(t)$.

The set of aritmetic statements are defined according to the previous definition with $t = q$. Then the following string of symbols are statements in arithmetics: $x = y$, $\neg \exists x \exists y\,(y = x + 1)$, etc.

For each statement, one can enumerate the number of different variable symbols appearing in the statement. A variable is called *free variable* in a statement if it does not appear right behind a quantifier. For example, the statement $\neg \exists x \exists y\,((y = x + 1) \vee (z = 1))$ has three variable symbols: x, y, and z. However, both the x and the y appears behind a quantifier. Hence, the only free variable of this statement is z.

8.2 Proofs

8.2.1 The Proof of Proposition 1

For all i, $d_q \in D_q$, and $m = (m^i, m^{-i})$ define

$$H_i^{d_q}(m) = \left\{ d_q\left(m^{i\prime}, m^{-i}\right) : m^{i\prime} \in \mathbb{N} \right\}.$$

Notice that $H_i^{d_q}(m) \subset 2^{A_q}$. Now, consider $\mathcal{H}_q : D_q \to 2^{A_q} \times \prod_{i=1}^{k} 2^{2^{A_q}}$ defined as

$$\mathcal{H}_q(d_q) = \left\{ \left(d_q(m), H_1^{d_q}(m), ..., H_k^{d_q}(m) \right) : m \in \mathbb{N}^k \right\}.$$

Then we have

Lemma 2 *Suppose that d, $d' \in D$ and $\mathcal{H}_q(d_q) = \mathcal{H}\left(d'_q\right)$ for all $q = 1, ..., n$. Then $G_d \sim G_{d'}$.*

We are ready to prove Proposition 1. For all $d_q \in D_q$, define

$$S(d_q) = \left\{ d_{-q} : \exists c_q \; \mathcal{H}_q\left(c_q\left([c_q], \left[c_{-q}^*\right]\right)\right) = \mathcal{H}_q(d_q), \; \mathcal{H}_q\left(c_{-q}^*\left([c_q], \left[c_{-q}^*\right]\right)\right) = \mathcal{H}_q(d_{-q}) \right\}.$$

By Lemma 2, it is enough to show that $\cap_{d_q \in D_q} S(d_q) \neq \{\emptyset\}$. Let us assume by contradiction that

$$\cap_{d_q \in D_q} S(d_q) = \{\emptyset\}.$$

Then there exists a function, $\mathcal{F} : \mathcal{H}_{-q}(D_{-q}) \to \mathcal{H}_q(D_q)$, such that,

$$d_{-q} \notin S(d_q) \tag{21}$$

if $\mathcal{F}(\mathcal{H}_{-q}(d_{-q})) = \mathcal{H}_q(d_q)$. Now, define a function $g : \mathcal{H}_q(D_q) \to D_q$, such that $\mathcal{H}_q(g(\mathcal{H}_q(d_q))) = \mathcal{H}_q(D_q)$. Furthermore, define the function $f : D_{-q} \to D_q$, such that $f(d_{-q}) = g(\mathcal{F}(\mathcal{H}_{-q}(d_{-q})))$. Notice that the domains of both $\mathcal{F}$ and g are finite, and hence, f is a definable function. In addition,

$$d_{-q} \notin S(f(d_{-q})), \tag{22}$$

by (21) and the definitions of $\mathcal{F}$ and g. Finally, we are ready to prove the proposition. Define the following contract for Principal q in one free variable:

$$c_q^x([c_q], [c_{-q}], m) = f\left(c_{-q}\left(\left[< x >^{(x)}\right], [c_{-q}], m\right)\right),$$

for all $m \in \mathbb{N}^k$. Let γ denote the Gödel code of this contract. Then

$$c_q^\gamma([c_q], [c_{-q}], m) = f\left(c_{-q}\left(\left[c_q^\gamma\right], [c_{-q}], m\right)\right). \tag{23}$$

First, notice that

$$c_{-q}\left(\left[c_q^\gamma\right], [c_{-q}]\right) \in S\left(c_q^\gamma\left(\left[c_q^\gamma\right], c_{-q}\right)\right) \tag{24}$$

by the definition of the set S. On the other hand, by (22),

$$c_{-q}\left(\left[c_q^\gamma\right], [c_{-q}]\right) \notin S\left(f\left(c_{-q}\left(\left[c_q^\gamma\right], [c_{-q}], m\right)\right)\right).$$

This can be rewritten by (23) as

$$c_{-q}\left(\left[c_q^\gamma\right], [c_{-q}]\right) \notin S\left(c_q^\gamma([c_q], [c_{-q}])\right).$$

But this contradicts to (24), and hence, $\cap_{d_q \in D_q} S(d_q) \neq \{\emptyset\}$.

8.2.2 Proof of Remark 2

Suppose that (d^*, β^*, α^*) implements an allocation in the ordinary contracting game. We construct equilibrium strategies in the contractible contracting game which implements the same allocation. Define the contract for Principal j, $c_j^{i*} \in C_j$, as follows: $c_j^{i*}\left(([c_q])_{q=1}^n, [m_j^i]\right) = d_j^{i*}\left([m_j^i]\right)$ for all $c \in C$ and $m_j^i \in \mathcal{L}$. In a subgame generated by the contract profile c, define the second-stage strategies of the agents as $\alpha^*(d(c), .)$, and the third stage strategies of the principals as $\beta^*(d(c), .)$.

We must show that the players have no incentives to deviate. First, recall that the subgame generated by c in the contractible contracting game is the same as the subgame generated by $d(c)$ in the ordinary contracting game. Since β^* and α^* were equilibrium strategies in the ordinary contracting game, all players play a Weak Perfect Bayesian Equilibrium in every subgame. We only have to show that principals have no incentive to deviate at the contracting stage. Suppose that Principal j offers the contract c_j instead of c_j^*. This deviation results a subgame generated by $\left(c_j\left([c_j], [c_{-j}^*]\right), d_{-j}^*\right)$. Notice that Principal j could generate the same subgame in the ordinary contracting game by offering $c_j^i\left([c_j], [c_{-j}^*], .\right)$ for all $i = 1, ..., k$. Since this deviation was not profitable in the ordinary contracting game, offering c_j is not profitable in the contractible contracting game.

8.2.3 Proof of Theorem 2

By Theorem 1, we only have to show that $E_t\mathcal{U}_j(t) = \underline{u}_j$ for all j. Consider first the following ordinary contract of Principal q $(q \neq j)$ offered to Agent i:

$$\widetilde{d}_q^i(l) = \left\{a_{q,j}^i\right\} \text{ for all } l \in \mathbb{N}.$$

Suppose that Principal q offers $\widetilde{d}_q = \left(\widetilde{d}_q^1, ..., \widetilde{d}_q^k\right)$ for all $q (\neq j)$. Then Principal $q (\neq j)$ ends up taking action $a_{q,j}$ regardless of what the messages of the agents and the contract of Principal j. Therefore, by part (ii) of Assumption 1, the expected payoff of Principal j is at most $\mathcal{U}_j(t)$ in every subgame $G_{\left(d_j, \widetilde{d}_{-j}\right)}$. Hence, $E_t\mathcal{U}_j(t) \geq \underline{u}_j$.

Now, consider the following contract of Principal j offered to Agent i:

$$\widetilde{d}_j^i(l) = \{\overline{a}_j\} \text{ for all } l \in \mathbb{N}.$$

Suppose that Principal j offers $\widetilde{d}_j = \left(\widetilde{d}_j^1, ..., \widetilde{d}_j^k\right)$. Then Principal j ends up taking action $\overline{a}_j$ no matter what the messages of the agents and the contracts of the other Principals are. Therefore, by part (i) of Assumption 1, the expected payoff of Principal j is at least $\mathcal{U}_j(t)$ in every subgame $G_{\left(\widetilde{d}_j, d_{-j}\right)}$. Hence, $E_t\mathcal{U}_j(t) \leq \underline{u}_j$.

8.3 Proof of Theorem 3

Since information is complete, every allocation is strongly incentive compatible. By Theorem 1, we only have to show that $\underline{u_j} = u_j^*$. Notice that Principal j can offer the contract d_j^i to Agent i, such that $d_j^i(l) = A_j^i$ for all $l \in \mathbb{N}$. This means that no matter what the messages of the agents are, Principal j can take any of his actions in the subgame generated by the contracts. Therefore, Principal j can best-respond to the action profile of the other principals and can achieve a value of at least u_j^*. This shows that $\underline{u_j} \leq u_j^*$. In order to prove that $\underline{u_j} \geq u_j^*$, let $a_{-j,j} = (a_{q,j})_{q \neq j} \in A_{-j}$ be a solution to $\min_{a_{-j}} \max_{a_j} u_j(a_j, a_{-j})$. Define Principal q's contract to Agent i, d_q^i, as $d_q^i(l) = a_{q,j}$ for all $l \in \mathbb{N}$. That is, no matter what the messages of the agents are, the principals other than j will take action $a_{-j,j}$. Of course, Principal j can achieve at most u_j^*, hence, $\underline{u_j} \geq u_j^*$.

References

[1] Arbatskaya, M., Hviid, M. and Shaffer, G. On the Incidence and Variety of Low-Price Guarantees. *The Journal of Law and Economics*, 47(1): 307-332, 2004.

[2] Arbatskaya, M., Hviid, M. and Shaffer, G. On the use of low-price guarantees to discourage price cutting. *International Journal of Industrial Organization*, 24(6):1139-1156, 2006.

[3] Aguzzoni, Luca et al.: Can 'Fair' Prices Be Unfair? A Review of Price Relationship Agreements, report of Office of Fair Trade, 2012.

[4] Andrea Attar, Thomas Mariotti, François Salanié. Nonexclusive Competition in the Market for Lemons. *Econometrica*, 79(6): 1869–1918.

[5] Kyle Bagwell and Robert W. Staiger. Reciprocity, non-discrimination and preferential agreements in the multilateral trading system. *European Journal of Political Economy*, 17(2):281–325, 2001.

[6] Baye, M. and D. Kovenock. How to Sell a Pickup Truck: 'Beat-or-Pay' Advertisements as Facilitating Devices *International Journal of Industrial Organization* 12, 21-33, 1994.

[7] Belton, T. A model of duopoly and meeting or beating competition. *International Journal of Industrial Organization* 5, 399–417, 1987.

[8] Chen, Z. How Low is a Guaranteed-Lowest-Price? *Canadian Journal of Economics* 28, 683-701, 1995.

[9] L. Epstein and M. Peters. A revelation principle for competing mechanisms. *Journal of Economic Theory*, 88(1):119–161, September 1999.

[10] C Fershtman and K.L. Judd. Equilibrium incentives in oligopoly. *American Economic Review*, 77:927–940, 1987.

[11] Hess, J. D. and Gerstner, E. Price-matching policies: An empirical case. *Managerial and Decision Economics*, 12(4): 305-315, 1991.

[12] Seungjin Han. Menu theorems for bilateral contracting. *Journal of Economic Theory*, 127(1):157–178, 2006.

[13] Adam Kalai, Ehud Kalai, Ehud Lehrer and Dov Samet. A Commitment Folk Theorem. *Games and Economic Behavior*, 69(1): 127-137, 2010.

[14] Michael Katz. Observable contracts as commitments: Interdependent contracts and moral hazard. *Journal of Economics and Management Strategy*, 15(3):685–706, 2006.

[15] Logan, J. and R. Lutter. Guaranteed Lowest Prices: Do They Facilitate Collusion. *Economics Letters* 31, 189-192, 1989.

[16] Mañez, J. A. Unbeatable Value Low-Price Guarantee: Collusive Mechanism or Advertising Strategy? *Journal of Economics & Management Strategy*, 15(1), pp. 143–166, 2006.

[17] David Martimort and Lars Stole. Communications spaces, equilibria sets and the revelation principle under common agency. University of Chicago unpublished manuscript, 1998.

[18] David Martimort and Lars Stole. The Revelation Principle and the Delegation Principle in Common Agency Games. *Econometrica*, 70(4): 1659-1673.

[19] Mas-Colell, A., M. Whinston and J. Green. *Microeconomic Theory.* New York, Oxford University Press, 1995.

[20] James Peck. A Note on Competing Mechanisms and the Revelation Principle. Ohio State University unpublished manuscript, 1997.

[21] Peters, M. Common agency and the revelation principle. *Econometrica*, 69(5):1349–1372, 2001.

[22] Peters M. and Szentes. Definable and Contractible Contracts. *Econometrica*, 80(1):363–411, 2012.

[23] Png, I.P.L. and D. Hirshleifer. Price Discrimination Through Offers to Match Price. *Journal of Business,* 60(2): 365-383, 1987.

[24] Rothschild, M. and J. E. Stiglitz. Equilibrium in Competitive Insurance Markets: An Essay in the Economics of Imperfect Information. *Quarterly Journal of Economics*, 87(3): 629-649.

[25] S.C. Salop. *Practises that Credibly Facilitate Ologopoly Coordination.* Cambridge, MIT Press, 1986.

[26] Leo Simon and W. Zame. Discontinuous games and endogenous sharing rules. *Econometrica*, 58(4):861–872, 1990.

[27] Moshe Tennenholtz. Program Equilibrium. *Games and Economic Behavior*, 49(2):363–373,2004.

[28] Takuro Yamashita. A revelation principle and a folk theorem without repetition in games with multiple principles and agents. *Econometrica*, 78(2):791-801, 2010.

[29] Zhang, J. Price-matching Policy and the Principle of Minimum Differentiation. *Journal of Industrial Economics* 43, 287-299, 1995.