Economic Theory (2007) 31: 539–552
DOI 10.1007/s00199-006-0111-9

RESEARCH ARTICLE

Levent Koçkesen

Unobservable contracts as precommitments

Received: 22 October 2005 / Accepted: 27 March 2006 / Published online: 22 April 2006
© Springer-Verlag 2006

Abstract It is well known that signing publicly *observable* contracts with third parties is a means of credibly committing to certain actions and hence may yield strategic advantages. Previous work on the commitment value of *unobservable* contracts has been limited to normal form games and extensive form games in which only one party has the option to sign a contract. In this paper, we extend the analysis to extensive form games in which both players can sign contracts, and characterize the set of sequential equilibria. We show that any Nash equilibrium outcome of the original game in which both players receive more than their individually rational payoffs can be supported as a sequential equilibrium outcome. Therefore, delegation acts not only as a commitment device to gain advantage over the opponent, but also as a cooperative device to attain Pareto improvements over the subgame perfect equilibrium outcome.

Keywords Strategic delegation · Unobservable contracts

JEL Classification Numbers C72 · D80 · L13

1 Introduction

It is well known that precommitment to certain actions in strategic interactions can yield a player payoff advantages. However, in many cases such commitments are ex

I would like to thank Atila Abdulkadiroglu, Jean-Pierre Benoit, Alberto Bisin, Boyan Jovanovic, Ehud Kalai, Giuseppe Lopomo, George Mailath, Efe Ok, Ariel Rubinstein, Andy Schotter, seminar participants at various universities and conferences, and anonymous referees for helpful comments and suggestions. Support from the Alfred P. Sloan Foundation and Program for Economic Research at Columbia University is also gratefully acknowledged.

L. Koçkesen
Department of Economics, Koç University,
Sariyer 34450, Istanbul, Turkey
E-mail: lkockesen@ku.edu.tr

post suboptimal and therefore not credible. The analysis of devices that make this type of precommitments credible has been one of the central themes in many areas of economics. As discussed in Schelling (1960) and demonstrated in many settings, contracts with third parties in general, and delegation in particular, may serve as one such commitment device. The observation that many interactions in the real world take place among agents makes this idea an intriguing one and brings about the question of whether delegation could take place due only to strategic reasons, i.e., as a means of credible precommitment. Under the assumption of observable contracts, previous literature has provided an affirmative answer to this question in many settings.[1]

If the contracts are unobservable, however, Katz (1991) showed that the Nash equilibrium outcomes of a game without delegation and those of the same game played between agents are identical. This seems to suggest that delegation via unobservable contracts has no effect on the expected outcome of a strategic situation and hence cannot yield any strategic advantage to the delegating party.

This is indeed an important observation in normal form games where the set of Nash equilibria may be regarded as the set of predicted outcomes. For example, in a standard Cournot duopoly game, Katz's result implies that the equilibrium outcome of the game is the standard Cournot outcome, irrespective of whether the game is played between principals or between agents. However, in extensive form games some type of sequential rationality is assumed on the part of the players and this makes the set of predicted outcomes generically smaller than the set of Nash equilibrium outcomes. The important question is, therefore, whether unobserved contracts have commitment value when the analysis is limited to, say, sequential equilibrium outcomes.[2]

Recently, Koçkesen and Ok (2004) addressed this question within the context of finite two-person extensive form games with a unique subgame perfect equilibrium outcome, called the *principals-only game* . In their framework, a *one-sided delegation game* is defined so that at the beginning of the game, one (and only one) of the principals decides whether to play the game himself or hire an agent by offering an incentive contract at a cost. If he chooses not to delegate, then the principals-only game is played by the two principals. If, on the other hand, he offers a contract, then the agent either accepts or rejects the contract. Upon rejection, two principals play the original game, whereas upon acceptance it is the agent and the other principal who play the original game. The crucial assumption is that the outside principal only observes whether he is playing against the agent or the principal, but not the contract.

Koçkesen and Ok (2004) show that the set of sequential equilibrium outcomes of the one-sided delegation game is equal to a subset of Nash equilibrium outcomes of the principals-only game in which (1) the outside principal behaves sequentially rationally and (2) the delegating principal receives a payoff greater than or equal to his subgame perfect equilibrium payoff in the principals-only game. In particular, in all equilibria of the delegation game in which the principal delegates, he receives

[1] See Brander and Spencer (1985), Eaton and Grossman (1986), Fershtman and Judd (1987) and Gatsios and Karp (1991), among others.

[2] Katz himself realizes the possibility of an affirmative answer to this question and demonstrates it via an example. Also, see Fershtman and Kalai (1997) for a demonstration of this possibility within the context of an ultimatum bargaining game.

a payoff *strictly* greater than his subgame perfect equilibrium in the principals-only game.

To illustrate how Koçkesen and Ok (2004) improve upon Katz (1991), consider a Stackelberg duopoly model with zero costs and inverse demand function

$$p(q_L + q_F) = \begin{cases} a - (q_L + q_F) & q_L + q_F \leqslant a \\ 0 & q_L + q_F > a \end{cases},$$

where q_L and q_F are the output choices of the leader and the follower, respectively. Katz's result, when applied to this game, tells us that the set of Nash equilibrium outcomes of the above game is the same as the set of Nash equilibrium outcomes of the delegation game induced by it, i.e., $\{(q_L, (a - q_L)/2) : q_L \leqslant a\}$. Since this is a sequential move game, however, this result is not very informative. A more interesting question is whether there exists a sequential equilibrium outcome of the delegation game that is different from the subgame perfect equilibrium (SPE) outcome of the principals-only game, which is $(a/2, a/4)$. Katz's analysis cannot answer this question: It only implies that the set of sequential equilibrium outcomes is a subset of the set of Nash equilibrium outcomes of the principals-only game, i.e., $\{(q_L, (a - q_L)/2) : q_L \leqslant a\}$. The results in Koçkesen and Ok (2004), however, imply that if only the leader has the option to delegate, then the answer to this question is negative, i.e., the unique sequential equilibrium outcome of the delegation game is $(a/2, a/4)$, in which the leader chooses not to delegate. If, on the other hand, only the follower has the option to delegate, then the set of sequential equilibrium outcomes of the delegation game is $\{(q_L, (a - q_L)/2) : q_L \leqslant a/2\}$, which includes only those Nash equilibrium outcomes that give the follower a (weakly) higher payoff than his SPE payoff in the principals-only game.

Unfortunately, the analysis of delegation games in Koçkesen and Ok (2004) remains incomplete, because it falls short of describing the set of possible outcomes that may arise in delegation environments in which both principals have the option to delegate. In the above Stackelberg game, for example, the results of Koçkesen and Ok (2004) do not characterize the set of equilibrium outcomes when both the leader and the follower have the option to delegate. Since delegation benefits the follower when he is the only player who has the option of delegating, a crucial question is whether giving the same option to the leader as well, counteracts this beneficial effect.

The present paper extends the analysis to delegation games where *both* principals have the option to delegate and characterizes the set of sequential equilibrium outcomes of *two-sided delegation games*. In the first stage of such a game principals decide between playing the game themselves or offering a contract to an agent. We assume that neither the delegation decision nor the contract choice is observable to the other principal-agent pair in this stage. In the second stage, all players learn whether they are facing a principal or his agent and the original (principals-only) game is played. Contracts remain unobservable throughout the game. We show that the possibility of the other principal delegating enlarges the set of sequential equilibrium outcomes significantly: the set of sequential equilibrium outcomes of the delegation game is *equal* to a certain subset of Nash equilibrium outcomes of the principals-only game. In particular, equilibrium outcomes of the one-sided delegation game remain as equilibrium outcomes, i.e., giving the option to delegate to the other player as well does not necessarily nullify the commitment advantage of delegation.

Furthermore, and perhaps more interestingly, any Nash equilibrium outcome of the principals-only game in which the principals receive more than their individually rational (minmax) payoffs can be supported as a sequential equilibrium outcome of the delegation game. This indicates the potential use of delegation not only as a commitment device to gain advantage over the opponent (for example in bargaining situations), but also as a cooperative device to attain Pareto improvements over the subgame perfect equilibrium outcome.

Recall that Katz's analysis would only imply that the set of sequential equilibrium outcomes of the delegation game is a *subset*of Nash equilibrium outcomes of the principals-only game, which includes the possibility that the unique sequential equilibrium outcome of the delegation game is the subgame perfect equilibrium outcome of the principals-only game. We show that, generically, this is not the case.

In the Stackelberg game described above, for example, the set of sequential equilibrium outcomes of the delegation game is $\{(q_L, (a - q_L)/2) : q_L \leq a\}$, which is exactly the set of Nash equilibrium outcomes of the principals-only game. Notice that in all the outcomes in this set the leader does at most as well as in the SPE outcome, whereas the follower may do better or worse. Therefore, in two-sided delegation games, the principals might do better or worse relative to the subgame perfect equilibrium outcome, whereas the delegating principal always benefits in one-sided delegation games.

The paper is organized as follows. Section 2 introduces some notation and presents the framework within which we shall conduct the analysis. Section 3 presents the main results and section 4 concludes with some remarks. All proofs are relegated to section 5.

2 Delegation games

The *principals-only game* is any finite, two player perfect information extensive form game with perfect recall, which we denote by Γ. We assume that this game has a unique subgame perfect equilibrium outcome and a finite set of Nash equilibrium outcomes. Denote the subgame perfect equilibrium payoff profile by $\Pi^{\text{SPE}} = \left(\Pi_1^{\text{SPE}}, \Pi_1^{\text{SPE}}\right)$, and the set of Nash equilibrium payoff profiles by $\Pi^{\text{NE}}(\Gamma)$. Let $NE_i^*(\Gamma)$ denote the set of Nash equilibria of Γ in which the behavior strategy of player i is sequentially rational after *every* history. We denote the set of payoff profiles that correspond to the strategy profiles in $NE_i^*(\Gamma)$ by $\Pi^{\text{NE}_i^*}(\Gamma)$. Let $\Pi_i(b)$ be the expected payoff of player i under behavior strategy profile b and define

$$\underline{\Pi}_1 \equiv \min_{b_2} \max_{b_1} \Pi_1(b) \text{ and } \underline{\Pi}_2 \equiv \min_{b_1} \max_{b_2} \Pi_2(b),$$

as the individually rational payoffs of player 1 and 2 in game Γ, respectively.

In a *delegation environment*, both principals have the option to play the game themselves or hire an agent via a contract. If principal $i \in \{1, 2\}$, offers a contract, the agent, denoted A_i, may accept or reject the offer. In this stage of the game (*delegation phase*) neither of the two principal-agent pairs is informed whether the other principal has decided to offer a contract and whether the agent has accepted it. We assume that if the agent rejects the contract, the game ends and the principal who has offered the contract receives a payoff of $-\infty$.[3] If the agent is offered a contract and she accepts it, then she plays the game in the second stage of the game (*game phase*). We assume that in the game phase every player knows the identity of the player she is facing without knowing the contract, if any, offered by the rival principal. One important assumption that we will maintain throughout the paper is that contracts cannot be renegotiated in the game phase. We discuss the renegotiation issue in some detail in section 4.

We assume that both agents' outside options are $\delta > 0$ and contracting costs $c > 0$ to both principals. Contracts condition the monetary payment to the agent on the outcome of the principals-only game and whether the agent plays the game against a principal or an agent. As a convention we denote the outcome z of the principals-only game obtained against a principal by $(z, 0)$ and against an agent by $(z, 1)$. For the results of this paper it is sufficient to limit the contract space to the simple set of contracts that pay either zero or the outside option to each outcome, i.e., the *contract space* is $\mathbb{C} \equiv \{0, \delta\}^{Z\times\{0,1\}}$, where Z is the set of outcomes in the principals-only game. Note that contracts can be conditioned only on the pure outcomes of the game rather than the agent's strategy and therefore are more realistic. However, they constrain the principal's ability to control his agent's actions in the game and introduce a number difficulties in the analysis.

The payoffs of a principal are not altered if he chooses not to delegate. If, on the other hand, a principal offers a contract, he incurs the cost c and pays the promised compensation to the agent in case the contract is accepted. The agent's payoff function is specified by the contract if she accepts it.

The description provided above induces a four-person extensive form game with imperfect information, $\Lambda(\Gamma, \delta, c)$, which we call a *delegation game,* and sometimes refer to as Λ.

3 The main result

Take any principals-only game Γ and consider the induced delegation game $\Lambda(\Gamma, \delta, c)$ as defined in the previous section. The first question that we are interested in is whether the set of equilibrium outcomes of $\Lambda(\Gamma, \delta, c)$ is different from that of the game Γ.[4] As we have mentioned previously, Katz (1991) has shown that the answer to this question is negative if one limits the analysis to Nash equilibrium outcomes. However, in extensive form games the relevant set of outcomes are those in which players behave sequentially rationally. In our framework, there is a unique such outcome of the principals-only game, i.e., the subgame perfect equilibrium outcome. The relevant comparison is therefore between the subgame perfect equilibrium outcome of the principals-only game and the set of sequential equilibrium outcomes of the delegation game. If the latter includes only the SPE outcome, then we can say that delegation has no effect on the outcome of the strategic interaction.

[3] Alternatively, and perhaps more naturally, we could assume that if the agent rejects an offer, then the principal plays the game. However, this assumption introduces some technical redundancies without changing any of the results, because in equilibrium no contract offer is rejected.

[4] It is obvious that any outcome in game $\Lambda(\Gamma, \delta, c)$ induces an outcome in game Γ. When we make statements about an outcome of $\Lambda(\Gamma, \delta, c)$ relative to that of Γ, we have in mind the outcome that $\Lambda(\Gamma, \delta, c)$ induces in Γ.

If, on the other hand, there are sequential equilibrium outcomes of $\Lambda(\Gamma, \delta, c)$ that are different from the subgame perfect equilibrium outcome of Γ, delegation can be said to affect the outcome and the question becomes understanding the nature of this effect. Our main result completely characterizes the set of sequential equilibrium outcomes of the delegation game in terms of the primitives of the principals-only game and therefore completely answers this question. It shows that for almost any principals-only game the set of sequential equilibrium outcomes of the delegation game includes those that are different from the subgame perfect equilibrium outcome of the principals-only game, and hence delegation may matter.

For purely technical and expositional reasons, we will restrict the analysis to a particular subset of sequential equilibria, in which

R1 The equilibrium strategy of the agent is to accept any contract offer that yields her an expected payoff greater than or equal to her outside option δ;
R2 Principals do not mix in their delegation decision in the sense that they either choose not to delegate or they choose a contract (or a probability distribution over the set of contracts) to offer.

We can show that all our results remain valid without R1, provided that we enlarge the contract space to $\mathbb{R}^{Z\times\{0,1\}}$ and define behavior strategies as simple probability measures. Relaxing restriction R2 changes the results only marginally without altering the main message of the paper, but introduces a significant expositional cost. We remark how the results change when this restriction is relaxed at the end of this section.

In order to state our main result we have to first define certain subsets of Nash equilibrium payoff profiles of the principals-only game. Let

$$\overline{\Pi}^{\mathrm{NE}_1^*}(\Gamma) \equiv \left\{(\Pi_1, \Pi_2) \in \Pi^{\mathrm{NE}_1^*}(\Gamma) : \Pi_2 > \Pi_2^{\mathrm{SPE}}\right\},$$
$$\overline{\Pi}^{\mathrm{NE}_2^*}(\Gamma) \equiv \left\{(\Pi_1, \Pi_2) \in \Pi^{\mathrm{NE}_2^*}(\Gamma) : \Pi_1 > \Pi_1^{\mathrm{SPE}}\right\},$$
$$\overline{\Pi}^{\mathrm{NE}}(\Gamma) \equiv \left\{(\Pi_1, \Pi_2) \in \Pi^{\mathrm{NE}}(\Gamma) : \Pi_1 > \underline{\Pi}_1, \Pi_2 > \underline{\Pi}_2\right\}.$$

In words, $\overline{\Pi}^{\mathrm{NE}_i^*}(\Gamma)$ is the set of Nash equilibrium payoff profiles of the principal-only game in which principal i plays sequentially rationally and principal $j \neq i$ receives a payoff strictly greater than his subgame perfect equilibrium payoff, whereas $\overline{\Pi}^{\mathrm{NE}}(\Gamma)$ is the set of Nash equilibrium payoff profiles in which both principals receive a payoff strictly greater than their respective minmax payoffs. Finally, let $\tilde{\Pi}^{\mathrm{SE}}(\Lambda)$ stand for the set of gross sequential equilibrium payoff profiles of the principals in the delegation game Λ. Then, we have the following result.

Theorem 1 *There exists an* $\ell > 0$, *such that*

$$\tilde{\Pi}^{\mathrm{SE}}(\Lambda) = \Pi^{\mathrm{SPE}} \cup \overline{\Pi}^{\mathrm{NE}_1^*}(\Gamma) \cup \overline{\Pi}^{\mathrm{NE}_2^*}(\Gamma) \cup \overline{\Pi}^{\mathrm{NE}}(\Gamma),$$

for all $\delta + c < \ell$.

Theorem 1 completely characterizes the set of sequential equilibrium payoffs of delegation games, and therefore improves upon the analysis in Katz, which would only imply that $\tilde{\Pi}^{SE}(\Lambda) \subseteq \Pi^{NE}(\Gamma)$. Furthermore, since an analogous statement to Theorem 1 also holds in terms of the outcomes, it also characterizes the set of sequential equilibrium outcomes. It demonstrates that delegation significantly enlarges the equilibrium outcome space of the principals-only game. It achieves this by freeing the principal from the requirement of sequential rationality, but not from rationality altogether. Furthermore, any Nash equilibrium outcome with payoffs greater than the individually rational (i.e., minmax) payoffs can be supported in equilibrium. This is possible because, through delegation, each principal can commit to punish the other if he fails to delegate.

The main idea behind the proof is fairly simple. To prove that any sequential equilibrium payoff profile must be in Π^{SPE}, $\overline{\Pi}^{NE_1^*}(\Gamma)$, $\overline{\Pi}^{NE_2^*}(\Gamma)$, or $\overline{\Pi}^{NE}(\Gamma)$ notice that there can be three types of sequential equilibria:

(1) Neither principal delegates. In this case sequential rationality of the principals implies that they must receive their subgame perfect equilibrium payoffs.
(2) Principal i delegates and principal $j \neq i$ does not delegate. The crucial observation in this case is that any contract offered in equilibrium must be such that the agent plays a best response (in terms of the principal's payoff function in the principals-only game) to the equilibrium strategy of the outside party in the game phase. If one of the principals does not delegate, then sequential rationality implies that his strategy must be not only a best response to the agent's strategy but it must also be sequentially rational. These two facts establish that sequential equilibrium payoffs must be in $\Pi^{NE_j^*}(\Gamma)$. The fact that delegating principal must receive a payoff strictly greater than his subgame equilibrium payoff follows from the fact that delegation is costly.
(3) Both principals delegate. In this case the agents' strategies must be best responses to each other in terms of the principals' preferences, but no sequential rationality is imposed. Therefore, the equilibrium gross payoffs must be a Nash equilibrium payoff profile. However, since either principal can guarantee the individually rational payoff by not delegating and since delegation is costly he must receive a payoff that is strictly greater than his minmax payoff.

The proof of the other direction constructs a sequential equilibrium for any payoff profile in Π^{SPE}, $\overline{\Pi}^{NE_1^*}(\Gamma)$, $\overline{\Pi}^{NE_2^*}(\Gamma)$, and $\overline{\Pi}^{NE}(\Gamma)$. Let f_δ be the contract that pays δ to any outcome. One can support a sequential equilibrium in which neither principal delegates and each receives his subgame perfect equilibrium payoff as follows. Principals play their respective subgame perfect equilibrium strategies in the game phase, whether they face a principal or an agent. Agent i plays principal i's subgame perfect equilibrium strategy after any history following the contract f_δ and beliefs put probability 1 on contract f_δ at all information sets. Notice that for principal i playing the subgame perfect equilibrium strategy is optimal even if he faces agent j in the game phase (which is an out of equilibrium event) because he believes that the agent has contract f_δ and plays the subgame perfect equilibrium strategy of principal j. The best deviation from equilibrium by a principal is to offer a contract which makes his agent to best respond to the rival principal's subgame perfect equilibrium strategy, which cannot yield a gross payoff that exceeds the subgame perfect equilibrium payoff.

The equilibrium in which principal 1 chooses not to delegate and principal 2 delegates (by offering the contract f_δ) and their payoffs are in $\overline{\Pi}^{NE_1^*}(\Gamma)$ can be supported as follows. Take a strategy profile (b_1, b_2) in $NE_1^*(\Gamma)$. Principal 1 plays b_1 if he faces agent 2 and agent 2 plays b_2 after contract f_δ. If principal 1 deviates and offers a contract, he believes that his agent will face an agent who plays b_2. Since b_1 is a best response to b_2, this cannot be a profitable deviation.

Finally, there are the equilibria in which both principals delegate and obtain their individually rational Nash equilibrium payoffs. These can be supported by making their agents play the Nash equilibrium strategy when facing the other agent and punish the other principal by minmaxing him if he fails to delegate.

The following two examples illustrate Theorem 1.

Example 1 (*Ultimatum bargaining*) Consider a simple ultimatum bargaining game in which player 1 gives either a low offer to player 2 (denoted l) or a high offer (denoted h), and player 2 either accepts (denoted y) or rejects (denoted n) the offer. If player 1 offers l and player 2 accepts, payoffs are 10 and 1 for player 1 and player 2, respectively. If player 1 offers h and player 2 accepts, then player 1 receives 1 and player 2 receives 10. If an offer is rejected, both players receive a payoff of zero. This game has two Nash equilibrium outcomes (l, y) and (h, y) among which only the outcome (l, y) is subgame perfect. Also notice that both (l, y) and (h, y) can be supported as NE_1^* equilibrium outcomes, whereas only (l, y) can be supported as a NE_2^* outcome. Therefore, we have the following payoff sets: $\Pi^{\mathrm{SPE}} = \{(10, 1)\}$, $\Pi^{\mathrm{NE}_1^*} = \{(10, 1), (1, 10)\}$, $\Pi^{\mathrm{NE}_2^*} = \{(10, 1)\}$, $\Pi^{\mathrm{NE}} = \{(10, 1), (1, 10)\}$, and $\left(\underline{\Pi}_1, \underline{\Pi}_2\right) = (0, 1)$. Furthermore, $\overline{\Pi}^{\mathrm{NE}_1^*} = \{(1, 10)\}$, $\overline{\Pi}^{\mathrm{NE}_2^*} = \emptyset$, and $\overline{\Pi}^{\mathrm{NE}} = \{(1, 10)\}$. Therefore, for small enough $\delta + c$, we have

$$\tilde{\Pi}^{\mathrm{SE}}(\Lambda) = \{(10, 1), (1, 10)\},$$

by Theorem 1. So, there are three types of equilibrium (1) neither principal delegates, (2) only principal 2 delegates, and (3) both delegate.[5] ‖

Example 2 (*Two-period alternating offers bargaining*) Consider now a slight modification of the game in Example 1. In the principals-only game, player 1 offers either l, m, or h and player 2 either accepts the offer or rejects it. If he accepts the offer the game ends and the payoff profile is $(10, 1)$ if the offer is l, $(5, 5)$ if the offer is m, and $(1, 10)$ if the offer is h. If player 2 rejects the offer, it is now his turn to offer l, m, or h. Player 1 either accepts or rejects the offer. If he accepts the offer and if the offer is l then the payoff profile is $(\theta, 10\theta)$, if the offer is m payoffs are $(5\theta, 5\theta)$, and if the offer is h the payoff profile is $(10\theta, \theta)$, where $\theta \in (1/2, 1)$ is the common discount factor. If player 1 rejects the offer, both players receive zero.

It is easily seen that the unique subgame perfect equilibrium outcome is for player 1 to offer h and player 2 to accept in the first round. Therefore, $\Pi^{\mathrm{SPE}} = \{(1, 10)\}$. One can also verify that $\left(\underline{\Pi}_1, \underline{\Pi}_2\right) = (\theta, 1)$, $\Pi^{\mathrm{NE}_1^*} = \{(1, 10), (\theta, 10\theta)\}$, $\Pi^{\mathrm{NE}_2^*} = \{(10, 1), (1, 10)\}$, and $\Pi^{\mathrm{NE}} = \{(1, 10), (10, 1), (5, 5), (5\theta, 5\theta), (\theta, 10\theta)\}$. Therefore,

$$\overline{\Pi}^{\mathrm{NE}_1^*} = \emptyset, \quad \overline{\Pi}^{\mathrm{NE}_2^*} = \{(10, 1)\}, \quad \overline{\Pi}^{\mathrm{NE}} = \{(1, 10), (5, 5), (5\theta, 5\theta)\},$$

[5] It can be shown that in this example $\tilde{\Pi}^{\mathrm{SE}}(\Lambda)$ can also be supported as a trembling hand perfect equilibrium payoff set. Therefore, Theorem 1 does not necessarily loose its validity if further refinements of sequential equilibrium is used.

and the set of gross sequential equilibrium payoff profiles is given by

$$\tilde{\Pi}^{\mathrm{SE}}(\Lambda) = \{(1, 10), (10, 1), (5, 5), (5\theta, 5\theta)\}.$$

Notice that since $\overline{\Pi}^{\mathrm{NE}_1^*}$ is empty there is no equilibrium in which only principal 2 delegates.

It is illustrative to use this example to emphasize the differences between the analyses in Katz (1991), Koçkesen and Ok (2004), and the present paper. Katz's result implies that

$$\tilde{\Pi}^{\mathrm{SE}}(\Lambda) \subseteq \Pi^{\mathrm{NE}}(\Gamma) = \{(1, 10), (10, 1), (5, 5), (5\theta, 5\theta), (\theta, 10\theta)\}.$$

In other words, his result is not informative enough to pin down the sequential equilibrium outcomes of the delegation game. The results in Koçkesen and Ok (2004), on the other hand, imply that if only player 1 has the option to delegate, then $\tilde{\Pi}^{\mathrm{SE}}(\Lambda) = \{(1, 10), (10, 1)\}$, and if only player 2 has the option to delegate, then $\tilde{\Pi}^{\mathrm{SE}}(\Lambda) = \{(1, 10)\}$, i.e., they precisely characterize the set of sequential equilibrium outcomes in environments in which only one of the principals has the option to delegate. In particular, in one-sided delegation games one can support only the most inequitable outcomes, i.e., outcomes (1, 10) and (10, 1). In contrast, the results of the present paper show that if both parties have the option to delegate, then delegation may yield outcomes in which players obtain more equitable outcomes, such as (5, 5).‖

Remark 1 The working paper version of the current paper Koçkesen (2004) relaxes restriction R2 and characterizes the complete set of sequential equilibrium outcomes of two-sided delegation games. It is shown that under certain conditions there are sequential equilibria in which the principals randomize between delegating and not delegating. In this case, certain convex combinations of payoffs in the sets $\Pi^{\mathrm{NE}}(\Gamma)$, $\Pi^{\mathrm{NE}_1^*}(\Gamma)$, and $\Pi^{\mathrm{NE}_2^*}(\Gamma)$ also arise as possible sequential equilibrium payoffs, in addition to the ones identified in Theorem 1.

4 Concluding remarks and extensions

In this paper we showed that the set of sequential equilibrium outcomes of a game played between agents is in general different from that of the same game played between the principals. In particular, virtually any Nash equilibrium outcome of the principals-only game can be supported as a sequential equilibrium outcome of the delegation game. Overall, our results suggest that contracts may still serve as effective precommitment devices even if they are unobservable and also help them attain Pareto superior outcomes. In the rest of this section, we discuss possible extensions and some of the limitations of the present framework.

4.1 Larger Classes of Principals-Only Games

Although we have limited our analysis to finite principals-only games with perfect information and a unique SPE outcome, it is easy to generalize Theorem 1 to larger classes of games. First, finiteness can be dispensed with, given that one defines an appropriate (sequential) equilibrium concept that can be applied to such games. Second, Theorem 1 remains valid in games with imperfect information with a unique SPE outcome. Third, although we have not attempted to do so, it seems straightforward to generalize Theorem 1 to games with non-unique SPE outcomes.

4.2 Forward Induction

In delegation games, there is always a sequential equilibrium in which the principals do not delegate and they obtain their SPE payoffs. However, Koçkesen and Ok (2004) show that, in one-sided delegation games, such equilibria are not well-supported, in the sense that beliefs or actions of the outside party do not conform to any equilibrium in which the principal actually chooses to delegate. Although the formalization of the well-supported equilibrium is different, the intuition behind this result can best be explained through a forward induction type argument: since delegation is costly, the outside party should interpret the fact that he is facing an agent, rather than the principal, as the principal being after an equilibrium that brings him a higher payoff than he would obtain by playing the game himself. Since, all such equilibria bring the principal a payoff strictly greater than his subgame perfect equilibrium payoff, the principal has a strict incentive to delegate if the cost of delegation is small enough.

In contrast, Koçkesen (2004) shows that in two-sided delegation games there may exist well-supported equilibria in which neither player delegates. In other words, the power of forward induction type refinements is weakened when both parties can delegate. However, it is shown that this occurs because in two-sided delegation games there may exist equilibria in which both principals completely mix between delegating and not delegating (see Remark 1). Therefore, if one limits the analysis to pure strategies, then in all well-supported equilibria at least one of the players chooses to delegate. Furthermore, even when mixed strategies are allowed, there is a large class of principals-only games in which delegation takes place with positive probability in *any* well-supported equilibrium of the delegation games that they induce. Koçkesen Koçkesen (2004) characterizes this class by a certain monotonicity condition regarding the way players' Nash equilibrium payoffs change as they commit to sequentially irrational strategies. Notably, it includes games that may be called 'competitive' in the sense that only one of the players strictly benefits from committing to a strategy that is not sequentially rational. Many interesting games, such as the ultimatum bargaining, entry-deterrence, and the chain-store games, belong to this class.

4.3 Renegotiation

An important question that is left unanswered in this paper is whether contracts have any commitment value if they can be renegotiated in the game phase.[6] In the present framework, frictionless renegotiation under symmetric information completely annihilates the commitment power of contracts, as it precludes the agent from acting sequentially irrationally from the perspective of the principal. Previous

[6] Renegotiation in the delegation phase, i.e., before the agents start playing the game, does not alter the results of this paper.

literature has identified two scenarios in which contracts may have commitment value even under renegotiation: (1) normal form games in which there is asymmetric information between the principal and the agent (see Caillaud et al. 1995); and (2) two-stage games with nontransferable utilities (see Bensaid and Gary-Bobo 1993). However, there are equally relevant environments in which renegotiable contracts may have commitment value. Indeed, Katz (1991) gives an example of an ultimatum bargaining game in which only the agent, but not the principal, observes the outside party's offer. He then constructs a sequential equilibrium in which a renegotiation proof contract yields the principal a payoff different from the unique SPE payoff of the principals-only game. In general extensive form games, it seems reasonable to conjecture that the set of equilibrium outcomes that can be achieved via renegotiable contracts depends on the extent and the nature of asymmetric information between the principal and the agent. In particular, depending upon the set of histories that are observed only by the agent, it could be possible to support outcomes other than the SPE outcomes of the principals-only game via renegotiable contracts. Although, such an inquiry is intriguing, it lies beyond the scope of the current paper and is left for future research.

5 Proofs

Consider a principals only game Γ and the induced delegation game $\Lambda(\Gamma, \delta, c)$. An *assessment* of Λ is a 2-tuple (β, μ), where $\beta = (\beta_i)_{i\in\{1,2,A_1,A_2\}}$ is a behavior strategy profile and μ a belief system. Let $SE(\Lambda)$ denote the set of sequential equilibria of the delegation game. We start by reminding the reader that in the set of sequential equilibria we are interested in, principals do not mix in their delegation decision (see R2 in section 3). In other words, $(\beta, \mu) \in SE(\Lambda)$ implies that each principal either plays the game himself or offers a contract (or a mixed strategy over the set of contracts).

Proof of Theorem 1 Define

$$\ell_1 \equiv \inf\left\{\Pi_1 - \Pi_1^{\text{SPE}} : \Pi_1 \in \Pi_1^{\text{NE}_2^*}, \Pi_1 > \Pi_1^{\text{SPE}}\right\},$$
$$\ell_2 \equiv \inf\left\{\Pi_2 - \Pi_2^{\text{SPE}} : \Pi_2 \in \Pi_2^{\text{NE}_1^*}, \Pi_2 > \Pi_2^{\text{SPE}}\right\},$$
$$\ell_3 \equiv \inf\left\{\Pi_1 - \underline{\Pi}_1 : \Pi_1 \in \Pi_1^{\text{NE}}, \Pi_1 > \underline{\Pi}_1\right\},$$
$$\ell_4 \equiv \inf\left\{\Pi_2 - \underline{\Pi}_2 : \Pi_2 \in \Pi_2^{\text{NE}}, \Pi_2 > \underline{\Pi}_2\right\},$$

and let $\ell \equiv \min\{\ell_1, \ell_2, \ell_3, \ell_4\}$. Clearly, $\ell > 0$.

We will first prove that $\tilde{\Pi}^{\text{SE}}(\Lambda) \subseteq \Pi^{\text{SPE}} \cup \overline{\Pi}^{\text{NE}_1^*}(\Gamma) \cup \overline{\Pi}^{\text{NE}_2^*}(\Gamma) \cup \overline{\Pi}^{\text{NE}}(\Gamma)$. The crucial observation used in the proof is that each contract offered with positive probability in equilibrium must be such that the agent best responds to the other party's strategy from the perspective of her principal's preferences. Let $(\beta, \mu) \in SE(\Lambda)$ and $\tilde{\Pi}_i$ be the gross equilibrium payoff of principal i. There are four cases to consider:

Case 1 Neither principal delegates. In this case principals play the game themselves and sequential rationality implies $\tilde{\Pi}_i = \Pi_i^{\text{SPE}}$.

Case 2 Principal 1 does not delegate and principal 2 delegates. Sequential rationality of principal 1 implies that his strategy must be a best response to agent 2's strategy after any history. Agent 2 on the other hand must be playing a strategy that is a best response (not necessarily sequentially rational) to principal 1's strategy from the perspective of principal 2's preferences. Therefore, the equilibrium outcome must be a $NE_1^*(\Gamma)$ outcome and we have $(\tilde{\Pi}_1, \tilde{\Pi}_2) \in \Pi^{NE_1^*}(\Gamma)$ (see Koçkesen and Ok (2004) for details of this argument). The fact that $\tilde{\Pi}_2 > \Pi_2^{SPE}$ follows from the fact that delegation is costly. Indeed, if $\tilde{\Pi}_2 \leqslant \Pi_2^{SPE}$, then (the net payoff) $\tilde{\Pi}_2 - \delta - c < \Pi_2^{SPE}$, and thus principal 2 could get a higher payoff by not delegating. Therefore, we have $\left(\tilde{\Pi}_1, \tilde{\Pi}_2\right) \in \overline{\Pi}^{NE_1^*}(\Gamma)$.
Case 3 Principal 1 delegates and principal 2 does not delegate. We can show, in a similar way to Case 2 above, that $(\tilde{\Pi}_1, \tilde{\Pi}_2) \in \overline{\Pi}^{NE_2^*}(\Gamma)$.
Case 4 Both principals delegate. Each agent must be best responding to the other agent's strategy from the perspective of her principal's preferences. Therefore, the outcome must be a Nash equilibrium outcome. Furthermore, the principals must be receiving a gross payoff that is strictly greater than their individually rational payoffs, since by deviating and playing the game themselves they would get at least as much and would not have to pay the cost of delegation. Therefore, we have $\left(\tilde{\Pi}_1, \tilde{\Pi}_2\right) \in \overline{\Pi}^{NE}(\Gamma)$.

To prove that $\tilde{\Pi}^{SE}(\Lambda) \supseteq \Pi^{SPE} \cup \overline{\Pi}^{NE_1^*}(\Gamma) \cup \overline{\Pi}^{NE_2^*}(\Gamma) \cup \overline{\Pi}^{NE}(\Gamma)$, define the contract f_δ as $f_\delta(z, \gamma) = \delta$ for all $(z, \gamma) \in Z \times \{0, 1\}$ and consider the following cases:[7]
Case 1 $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) = \Pi^{SPE}$. Let $\left(b_1^{SPE}, b_2^{SPE}\right)$ be a subgame perfect equilibrium of the principals-only game and consider the following assessment. Principal 1 and 2 choose not to delegate and they play their respective subgame perfect equilibrium strategies in the game phase. Agent i plays b_i^{SPE} after any history following the contract f_δ. Specify the strategies of the agents as any sequentially rational strategy following any contract other than f_δ and let the beliefs put probability 1 on contract f_δ at all information sets.

Under this assessment, equilibrium payoffs of the principals are their respective subgame perfect equilibrium payoffs. The best deviation from equilibrium by a principal is to offer a contract which makes his agent to best respond to the rival principal's subgame perfect equilibrium strategy. However, this cannot yield a gross payoff that exceeds the subgame perfect equilibrium payoff and hence there is no profitable deviation. It can be easily checked that strategies are sequentially rational given beliefs after any other history as well. Consistency of beliefs can be established by constructing a sequence of completely mixed assessments, in which contract f_δ receives probability $\varepsilon \in (0, 1)$ whereas any other contract receives a probability ε^K, where $K = |h_{\max}| + 2$ and $|h_{\max}|$ is the length of the longest history in Γ. All other non-equilibrium actions receive probability ε. As ε converges to zero, this sequence converges to the assessment described above.

[7] In other words the contract f_δ pays δ to any outcome. This contract makes it easy to satisfy the agent's sequential rationality and hence simplifies the proof. However, the result can be proved by using more reasonable contracts.

This assessment, therefore, is a sequential equilibrium of the delegation game with equilibrium payoffs given by the SPE payoffs.

Case 2 $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) \in \overline{\Pi}^{NE_1^*}(\Gamma)$. There exists, by hypothesis, a behavior strategy profile $(b_1, b_2) \in NE_1^*(\Gamma)$, such that $\Pi_i(b) = \hat{\Pi}_i$ and $\hat{\Pi}_2 > \Pi_2^{SPE}$. Consider the following assessment. Principal 1 chooses not to delegate and principal 2 delegates by offering the contract f_δ. Principal 1 plays b_1^{SPE} when facing the other principal in the game phase and plays b_1 when facing the agent 2. Principal 2 always plays b_2^{SPE} in the game phase. After any history following the contract f_δ, agent 1 plays b_1^{SPE} when facing principal 2 and b_1 when facing agent 2. Agent 2 plays b_2 after any history following the contract f_δ. Specify the strategy of agent i at other histories as any sequentially rational strategy and let the beliefs put probability 1 on contract f_δ at all information sets.

Clearly, the gross payoff profile of this assessment is $\left(\hat{\Pi}_1, \hat{\Pi}_2\right)$ and the net payoffs are given by $\left(\hat{\Pi}_1, \hat{\Pi}_2 - \delta - c\right)$. If principal 2 deviates by not delegating he obtains $\Pi_2^{SPE} < \hat{\Pi}_2 - \delta - c$, for $\delta + c < l$.[8] If he deviates by offering another contract, his agent will face principal 1, who plays b_1. Since b_2 is a best response to b_1, the best that he can get by such deviations is $\hat{\Pi}_2 - \delta - c$. These observations imply that there is no profitable deviation opportunities for principal 2. If principal 1 deviates and offers a contract, his agent will face another agent who plays b_2. Since b_1 is a best response to b_2, the most he can get is $\hat{\Pi}_1 - \delta - c < \hat{\Pi}_1$. Therefore, his strategy of not delegating is optimal as well. Sequential rationality at other histories and consistency are easily established. The proof of consistency is similar to Case 1 above.

Case 3 $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) \in \overline{\Pi}^{NE_2^*}(\Gamma)$. The proof of this case is symmetrical to that of Case 2 and, therefore, omitted.

Case 4 $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) \in \overline{\Pi}^{NE}$. In this case, there exists $(b_1, b_2) \in NE(\Gamma)$ such that $\Pi_i(b) = \hat{\Pi}_i$, $\hat{\Pi}_i > \underline{\Pi}_i$, $i = 1, 2$. Define $\Pi_i(b|h)$ as the expected payoff of principal i in Γ under strategy profile b, conditional upon reaching history h. Let

$$b_1^m \in \arg\min_{b_1}\{\max_{b_2} \Pi_2(b_1, b_2)\},\ b_2^m \in \arg\min_{b_2}\{\max_{b_1} \Pi_1(b_1, b_2)\},$$

be the minmax strategies and

$$b_1' \in \arg\max_{b_1} \Pi_1(b_1, b_2^m|h),\ b_2' \in \arg\max_{b_2} \Pi_2(b_1^m, b_2|h),\ \text{for all } h$$

be the principals' sequential best responses to the minmax strategies. Consider the following assessment. Both principals delegate by offering the contract f_δ and play their respective SPE strategies when they face the other principal in the game phase. Following the contract f_δ, agent i plays b_i^m when facing the other principal (i.e., minmaxes principal j)[9] and plays b_i (the Nash equilibrium strategy) when

[8] This is because $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) \in \overline{\Pi}^{NE_1^*}(\Gamma)$ implies $\hat{\Pi}_2 > \Pi_2^{SPE}$ by definition.

[9] Note that on the equilibrium path agent i plays against agent j, not principal j. Therefore, the fact that, under the equilibrium contract, agent i minmaxes principal j is an out of equilibrium event and does not contradict principal i's sequential rationality.

facing the other agent. Principal i plays a sequential best response to b_j^m when facing agent j. Specify β_{A_i} at other histories as any sequentially rational strategy, and let beliefs put probability 1 on contract f_δ at all information sets.

The gross equilibrium payoff profile of this assessment is $\left(\hat{\Pi}_1, \hat{\Pi}_2\right)$ and the net payoffs are given by $\left(\hat{\Pi}_1 - \delta - c, \hat{\Pi}_2 - \delta - c\right)$. If principal 1 deviates by not delegating he faces agent 2 playing b_2^m. The most he can get by such a deviation is $\max_{b_1} \Pi_1(b_1, b_2^m) = \underline{\Pi}_1 < \hat{\Pi}_1 - \delta - c$, for $\delta + c < l$.[10] Therefore, he has no profitable deviation opportunities. Similarly, there is no profitable deviation for principal 2 either. Sequential rationality at other histories and consistency of the beliefs are easily established. This completes the proof of Theorem 1. □

References

Bensaid, B., Gary-Bobo, R.J.: Commitment value of contracts under renegotiation constraints. Econometrica **61**, 1423–1429 (1993)

Brander, J.A., Spencer, B.J.: Export subsidies and international market share rivalry. J Int Econ **18**, 83–100 (1985)

Caillaud, B., Jullien, B., Picard, P.: Competing vertical structures: precommitment and renegotiation. Econometrica **63**, 621–646 (1995)

Eaton, J., Grossman, G.: Optimal trade and industrial policy under oligopoly. Q J Econ **101**, 383–406 (1986)

Fershtman, C., Judd, K.: Equilibrium incentives in oligopoly. Am Econ Rev **77**, 927–940 (1987)

Fershtman, C., Kalai, E.: Unobserved delegation. Int Econ Rev **38**, 763–774 (1997)

Gatsios, K., Karp, L.: Delegation games in custom unions. Rev Econ Stud **58**, 391–398 (1991)

Katz, M.: Game-Playing Agents: Unobservable contracts as precommitments. Rand J Econ **22**, 307–328 (1991)

Koçkesen, L.: Unobservable contracts as precommitments. Mimeo: Columbia University (2004)

Koçkesen, L., Ok, E. A.: Strategic delegation by unobservable incentive contracts. Rev Econ Stud **71**, 397–424 (2004)

Schelling, T: The strategy of conflict. Cambridge: Harvard University Press 1960

[10] This is because $\left(\hat{\Pi}_1, \hat{\Pi}_2\right) \in \overline{\Pi}^{\mathrm{NE}}(\Gamma)$ implies $\hat{\Pi}_i > \underline{\Pi}_i$, $i = 1, 2$.