Contract Complexity

Moshe Babaioff
Microsoft Research
Mountain View, CA 94043
moshe@microsoft.com

Eyal Winter*
Department of Economics
The Hebrew University
Jerusalem, Israel 91905
mseyal@mscc.huji.ac.il

April 29, 2014

Abstract

We study the complexity required for the implementation of multi-agent contracts under a variety of solution concepts. A contract is a mapping from strategy profiles to outcomes. Practical implementation of a contract requires it to be "simple", an illusive concept that needs to be formalized. A major source of complexity is the burden involving verifying the contract fulfillment (for example in a court of law). Contracts which specify a small number of outcomes are easier to verify and are less prone to disputes. We therefore measure the complexity of a contract by the number of outcomes it specifies. Our approach is general in the sense that all strategic interaction represented by a normal form game are allowed. The class of solution concepts we consider is rather exhaustive and includes Nash equilibrium with both pure and mixed strategies, dominant strategy implementation, iterative elimination of dominated strategies and strong equilibria.

Some interesting insights can be gained from our analysis: Firstly, our results indicate that the complexity of implementation is independent of the size of the strategy spaces of the players but for some solution concepts (but not all) grows with the number of players. Second, the complexity of *unique* implementation is sometimes slightly larger, but not much larger than non-unique implementation. Finally and maybe surprisingly, for most solution concepts implementation with optimal cost usually does not require higher complexity than the complexity necessary for implementation at all.

keywords: Contract Theory, Complexity.

JEL codes: C72, B86.

*Eyal Winter acknowledges the German-Israeli Science Foundation for its financial support.

1 Introduction

The practical implementation of contracts critically depends on their complexity. ”Simple” contracts require both signatories and the arbitrator to provide less evidence for the purpose of the contract verification. Simple contracts reduce the cost of implementation and increase the signatories’ trust in the contract effectiveness. However, since Holmstrom [1982] the literature studying the simplicity in contracting environments has been rather limited. This is probably so as the notion of ”simple” is rather illusive. In this paper we wish to study this issue in a general framework of multi-agent contracting. Our main objective is to study the tradeoff between the complexity needed to implement an arbitrary outcome through a contract, on one hand, and the equilibrium notion applied, on the other hand. Moreover, we would like to understand the tradeoff between complexity and the total payments made to the agents.

Our framework will be rather abstract and will lack a specific context: A principal who wishes to influence agents to play a specific profile of actions can use positive transfers to affect the structure of incentives. The exogenous strategic environment is given by an n-person normal form game and the principal’s transfers affect preferences in a qausi-linear manner. The contract maps conditions regarding the profile played to a vector of positive transfers to the agents. There is clearly more than a single way to define the complexity of a contract. For example, one might consider the computational complexity of the contract, or its Kolmogorov complexity [1998]. We focus on the ”evidence complexity” as explained below.

Generally, a contract might specify for every profile played, a different vector of transfers. Such a contract seems rather complex in the following sense. If the signatories would like to argue that a specific outcome (vector of payments) should have been made, they need to bring evidence regarding the profile played. Clearly, the same piece of evidence cannot be used to argue for different outcomes. Thus, the number of different pieces of evidence that are necessary and sufficient to argue about every profile played should be equal to the number of possible outputs of the contract. We measure the complexity of a contract by the number of possible outputs of the contract. Note that as the mapping between profiles played and outcomes is deterministic, such a mapping induces a partition over the set of strategy profiles, and our complexity measure counts the number of elements in the partition, thus we call it the *partition complexity* of the contract.

To further explain the motivation behind our definition of complexity we first note that the more refined the contract partition is, the less evidence a signatory is required to bring in order to verify the contract’s fulfillment (vis a vis himself or the court). For example assume that the principal contracts with two workers to jointly perform a project on his behalf with each agent being responsible for one task, and the two task are similar. A contract that specifies payment only as a function of the number of tasks performed (i.e., 0, 1, or 2) is properly simpler than a contract that also makes contingencies on the identity of the agent whose task ended successfully. The contract partition that corresponds to the latter contract is a refinement of the partition of the first one. Furthermore, the latter contract is more complex than the former. If only one agent completed his task successfully one would need to bring evidence regarding who this agent was and who failed, but with the former contract it is enough to bring evidence proving that only one agent completed his task successfully and no evidence concerning the person’s identity is needed.

While coarser contract partitions are simpler they are more limited in terms of the incentives that they set up for agents to play the strategy profile that the principal desires them to play. This is precisely the tradeoff that we wish to investigate in this paper; namely, the tradeoff between the simplicity of contracts and what they can achieve in terms of incentives.

Not every two contracts are compared by the simplicity criterion which is based on refining and coarsening of partitions. However, one can consider the number of components in the contract partition as a proxy for the refinement comparison. It is the most natural extension of the partial order induced by the refinement relation, to a complete order on the entire class of partitions. In most of our results when we compare contracts for their simplicity these contracts will be comparable by means of the partial order of refinement. In the exceptional cases in which these partitions are incomparable we will compare them by means of the complete order extension (i.e., the number of components of the partitions).

Our Results: We investigate how the complexity of the contract relates to the solution concept the principal is interested in implementing. Our analysis covers almost all the well known non-cooperative solution concepts. This includes: Nash equilibrium in both pure and mixed strategies, dominant strategy implementation, iterative elimination of dominated strategies, (non-equilibrium) undominated strategies and strong equilibrium. We will be referring to the last concept as a coalitional solution and to the rest as individual solutions.[1]

For all these solution concepts we will be concerned with both *unique implementation* (in which the principal's desired outcome is sustained as a unique equilibrium outcome) and non-unique implementation. Finally, for each solution concept and for each type of implementation (unique or non-unique) we will be concerned with both feasible implementation as well as optimal implementation. In a *feasible implementation* the principal is not concerned about the payments that he transfers to agent to sustain his desirable outcome. In contrast, in an *optimal implementation* the principal is constrained not to pay more than he would have if there was no constraint on the complexity of contracts. Put differently, the principal is in pursue of the simplest contract whose cost is identical to the cost of the contract in which he is allowed to make contingencies on each strategy profile. Clearly, feasible implementation requires (weakly) less complexity than optimal implementation and controlling for these two types of implementations, unique implementation require (weakly) more complex contracts than non-unique implementation.

Roughly, our results can be summarized as follows. For weak individual solution solution concepts (like Nash and undominated strategies), the complexity of optimal implementation is only 2, as simple as possible. Implementation as a unique pure Nash equilibrium requires slight higher complexity of 3, and this complexity is sufficient for optimal implementation. Once moving to more demanding solution concepts, like iterative elimination of dominated strategies, the complexity of feasible implementation grows to become linear in the number of agents (complexity of $n + 1$ when n is the number of agents). The same complexity is sufficient for dominant strategy implementation, but not for optimal implementation - such implementation has strictly higher complexity, but not much higher (still linear). For strong Nash equilibrium, a coalitional solution concept, we show that optimal implementation is simple, a complexity of 2 is sufficient. Such a low complexity is also sufficing for feasible implementation of unique strong Nash equilibrium. We leave open the complexity of an

[1]The distinction concerns on the type of deviation on which the equilibrium concept is built i.e. coalitional deviations vs. individual deviations.

optimal implementation of a profile as a unique strong Nash equilibrium.

Let us discuss some insights from the results. First our results indicate that the complexity of implementation is independent of the size of the strategy spaces of the players but for some solution concepts it grows (linearly) with the number of players. Interestingly, the contracts we present have the property that the principal needs to know very little about the game even if he would like to implement a profile in dominant strategies with optimal cost. All he needs to know is a *bound* on the payoffs of the game, and the payments needed for Nash implementation. All other details, like the exact space of strategies and the exact payoffs, are not needed. This means that unforeseeable actions that do not carry substantial payoffs to the parties cannot disrupt the effectiveness of the contracts. Additionally, the incompleteness of these contracts under such unforeseeable actions does not affect their optimality. Thus, these contracts are quite robust, and can be used even if the actually game is not exactly the one the principal believe the agents are playing.[2]

Second, our results indicate that insisting on *unique* implementation does not increase the complexity substantially. For individual solution concepts this is the case for both feasible and optimal implementation. The only exception is pure Nash equilibrium, but the increase in complexity is minor. For (mixed) Nash equilibrium the question remains open. For strong Nash equilibrium the uniqueness requirement does not imply an increase in complexity for feasible implementation (and the question for optimal implementation remains open).

Finally, our results provide important insight regarding the comparison between feasible and optimal implementation. Apriori one might expect a trade-off between the cost of the contract (the total payment awarded to agents) and its complexity. One might expect that reducing the complexity of a contract might come at the expense of higher power incentives offered to agents. Surprisingly this is generally not the case for individual solution concepts. Our findings do not indicate such a tradeoff, except of a slight increase in complexity for optimal implementation in dominant strategies. In fact, for most solution concepts and for both feasible and optimal implementations, a contract with minimal complexity for feasible implementation also achieves the first best, and hence also provides an optimal implementation. i.e., it's cost does not exceed the cost of a contract in which the principal can verify any strategy profile. Exceptional in this respect is the coalitional solution of unique strong Nash equilibria for which we haven't be able to establish this comparison though we conjecture that some tradeoff between complexity and cost does exist.

Related Work: The complexity of contracts has been discussed in the economics literature mainly in the context of bilateral contracting. Much of this literature builds on the important insight regarding incomplete contracts that was introduced in Williamson's [1975, 1985] two seminal books and later by Hart and Moore [1988]. According to this insight, contracts are never complete because of the enormous complexity of all possible contingencies. Holmstrom and Milgrom [1987] argue that "price-only" contracts are often used in the real world to avoid the complexity of efficient contracts.

Anderlini and Felli [1994] consider "algorithmic" contracts i.e. contracts for which the mapping from the states of the world to outcome can be computed in a finite number of steps. Although not every contract satisfies this condition, they are able to show that every contract can be approximated (in terms of signatories expected payoffs) by algorithmic contracts. Melamad et al [1997] discuss the relation between contract complexity and delegation in a

[2]Additionally, computing the contract is easy.

principal and 2-agent effort game. They show that delegation can reduce the complexity of contracts. They define complexity in terms of the number of contingencies. Dye [1985] considers context dependent complexity of contracts by introducing an explicit measure of complexity. In contrast, Anderlini and Felli [1999] use complexity cost to explain contractual incompleteness without referring to a specific measure of complexity. Instead they consider the allocation of property rights among two agents and establish results that quantify over a class of complexity cost functions to assess the tradeoff between complexity and efficiency. Segal [1999] considers contracts complexity in the context of the holdup problem. Complexity here is defined in terms of the number of future trading opportunities that are put into the contracts.

Unlike the above literature our model and analysis attempt to capture a general contracting environment based on an arbitrary n-person normal form game and a plethora of equilibrium solution concepts. While the tradeoff between complexity and efficiency has been discussed in the existing literature for various more specific environments, as far as we know, the tradeoff between complexity and the strength of the solution concept has not been studied before. One of our main objectives in this paper would be to find the least complex contract that allows an arbitrary outcome to be implemented for each such solution concept.

Recently, Hart and Nisan [2013] have studied the menu complexity of auctions. A seller aiming to maximize his revenue offers a menu to a buyer with a combinatorial valuation, and the complexity of a menu is the number of different elements in the menu. Note that much like our complexity measure, this measure also counts the number of different outcomes of the mechanism. The performance of simple mechanisms (with only few outcomes) was recently further studied by Li and Yao [2013] and Wang and Tang [2013].

Organization: In section 2 we set up the model list the definitions of the solution concepts we are dealing with and present the well known example of team production with two levels of effort, a game that will be useful for many of our results. In Section 3 we provide our analysis for individual solution concepts. Our analysis for the coalitional solution concepts is provided in Section 4. We conclude with comments in Section 5.

2 Model and Preliminaries

We are given a normal-form game $G = (N, S, U)$ with a set N of n agents, space of strategies S and utilities U. Agent i's set of strategies (actions) is S_i. The space S of strategy profiles is $S = S_1 \times \ldots \times S_n$.[3] When the agents play the strategy profile $s \in S$, agent i's utility in the game is $U_i(s)$. A principal is interested in implementing some arbitrary strategy profile $s^* \in S$ via some solution concept E.

The principal is able to sign a contract with the agents. A contract can only specifying non-negative payments from the principal to the agents (his inability to fine the agents is known as the *limited liability* constraint). Formally, a contract C specifies for each profile $s \in S$ and each agent $i \in N$ a payment $C_i(s) \geq 0$ to the agent. Given the game G and contract C, the induced game $G_C = (N, S, U + C)$ is a game with the same set of agents N and same strategy space S, but the payoff for agent i when the agents are playing strategy $s \in S$ is now $U_i(s) + C_i(s)$. This is so as agent i is also being paid $C_i(s)$, on top of his utility in the game G which was $U_i(s)$.

[3]Note that representing an arbitrary game requires exponential space.

Given a normal form game G, a *solution concept* E is defined by a set of strategy profiles $E(G) \subseteq S$. We say that a *contract* C *implements the profile* s^* *via the solution concept* E if s^* belongs to $E(G_C)$. We say that C implements s^* as the *unique* profile satisfying the solution concept E if $E(G_C) = \{s^*\}$.

We are interested in understanding the minimal complexity of contracts that implement an arbitrary desired profile s^* via various solution concepts. We consider the following way to represent a contract. A contract C is defined by a partition $\mathcal{P}$ of the space of strategies S to a disjoint union of k pieces $(P_1, P_2, \ldots, P_k)$, and for each element of the partition (part) the contract specifies the payment to each agent if the strategy profile belongs to that part: Formally, for each part P_j in the partition and every $s \in P_j$, the payment $C_i(s)$ to agent i is $C_i(P_j)$. We view the partition $\mathcal{P}$ as a constraint imposed on the contract in terms of contingencies that are allowed. The contract can make contingencies only on elements of the partition and is not allowed to distinguish between two strategy profiles belonging to the same element of the partition. The *partition complexity* of a contract is defined as the number of parts in the partition $\mathcal{P}$ of that contract. For a contract with partition $(P_1, P_2, \ldots, P_k)$ the complexity is k.

The *cost* of implementing profile s^* for a given contract C is the sum of payments to all agents, as determined by C, when the agents play the profile s^*, that is $\sum_{i \in N} C_i(s^*)$.[4] Given a solution concept E, the *optimal cost* of implementing profile s^* via E is the infimum cost of implementing s^* via E, by a contract C that can assign an arbitrary non-negative payment to each profile $s \in S$.

For a given $\epsilon > 0$, we say that a contract C implements the profile s^* via E with *ϵ-optimal cost*, if the cost of implementing the profile s^* via E by the contract C is larger than the optimal cost of implementing profile s^* via E by at most ϵ.

We are interested both in the minimal complexity needed to implement a profile via some solution concept, and in the minimal complexity needed to implement a profile via some solution concept *with optimal cost*. For a given solution concept E, we say that *feasible implementation in the solution concept E has partition complexity k* if for any game and any profile s^*, it is possible to implement the profile s^* via E by a contract of partition complexity k, but not by a contract of complexity $k-1$ (that is, for some game and for some profile s^* there exist no contract of complexity $k-1$ implementing s^*). Additionally, we say that *optimal-cost implementation in the solution concept E has partition complexity k* if for any game, any profile s^* and any $\epsilon > 0$, it is possible to implement the profile s^* via E with ϵ-optimal cost by a contract of partition complexity k, but not by a contract of complexity $k-1$.

2.1 Solution Concepts

We next formally define the solution concepts we consider, all of them are standard solution concepts and we present them for completeness.

Definition 1. Given a game G:

[4] We remark that the contract is allowed to promise larger payments if s^* is not played, these promises will not affect the cost of implementing s^* as long as they are not realized when s^* is played. We later observe that for dominant strategy implementation disallowing such promises might result in exponential size contracts.

- a profile of strategies $s^* \in S$ is a *pure Nash equilibrium* if for every $i \in N$ it holds that $U_i(s^*) \geq U_i(s_i, s^*_{-i})$ for every $s_i \in S_i$.

- $s^*_i \in S_i$ is a *strict dominant strategy* for $i \in N$ if $U_i(s^*_i, s_{-i}) > U_i(s_i, s_{-i})$ for every $s \in S$. A profile s^* is a (strict) dominant strategies profile if for every $i \in N$, the strategy s^*_i is a strict dominant strategy.

- $s^*_i \in S_i$ *dominates* strategy $s_i \in S_i$ if $U_i(s^*_i, s_{-i}) \geq U_i(s_i, s_{-i})$ for every $s_{-i} \in S_{-i}$, with strict inequality for some s_{-i}. A profile s^* is *undominated* if for every $i \in N$, the strategy s^*_i is not dominated by any other strategy.

- $s^* \in S$ survives *iterative removal of strictly dominated strategies* if there exist a sequence of strategies $s^1_{k_1}, s^2_{k_2}, \ldots, s^r_{k_r}$ that does not include the strategy s^*_i for every $i \in [n]$, such that for every $j \in [r-1]$ if strategies $s^1_{k_1}, s^2_{k_2}, \ldots, s^j_{k_j}$ are eliminated then $s^{j+1}_{k_{j+1}}$ becomes a dominated strategy for agent k_j, and there is not strategy that is not on the list that can be added as the $r+1$ element of the list.

- a profile of mixed strategies $s^* \in \Delta(S)$ is a *mixed Nash equilibrium* if for every $i \in N$ it holds that $U_i(s^*) \geq U_i(s_i, s^*_{-i})$ for every $s_i \in \Delta(S_i)$.

- a profile of strategies $s^* \in S$ is a *strong Nash equilibrium* if there is no coalition $R \subseteq N$ and strategy profile $s_R \in S_R$ such that for every $i \in R$ it holds that $U_i(s_R, s^*_{-R}) > U_i(s^*)$.

2.2 The Effort Game

One specific game that would be useful in proving lower bounds on the partition complexity of various solution concepts is the *Effort Game*.[5] In this game each agent $i \in N$ has a set of two possible actions $S_i = \{0, 1\}$, the low effort action ($s_i = 0$) has a cost of 0 independent of the actions of others, while the high effort action ($s_i = 1$) has a cost of $c_i > 0$ again, independent of the actions of others. In this game the principal will be interested in implementing the all-effort vector ($s_i = 1$ for all i). Note that if payments are not contingent at all on the strategies played, every player has a dominant strategy to shirk. This means that any non-trivial solution concept will require partition complexity of at least 2.

3 Individual Deviation Concepts

In this section we consider individual deviation concepts: Nash equilibrium, unique Nash equilibrium, iterative removal of strictly dominated strategies, and strict dominant strategies.

The *Nash cost* of implementing s^* is defined to be $\sum_{i \in N} \hat{t}^{NE}_i$ for

$$\hat{t}^{NE}_i = \max_{\hat{s}_i \in S_i} U_i(\hat{s}_i, s^*_{-i}) - U_i(s^*)$$

Clearly any implementation in a solution concept that is stronger than Nash equilibrium requires paying at least the Nash cost. We observe below that the Nash cost is the optimal cost not only to implement a profile as a Nash equilibrium, but also to implement it as

[5] A moral hazard version of this game was studied in [16, 3].

the unique Nash equilibrium, as the unique profile that survives iterative removal of strictly dominated strategies, or in strict dominant strategies.

3.1 (Non-unique) Nash Equilibrium

As mentioned, for the effort game one cannot hope to implement the all-effort vector as a Nash equilibrium with partition complexity of 1, as shirking is a dominant strategy for such a contract. One can immediately observe that optimal-cost implementation in Nash equilibrium has partition complexity 2. The partition that implements a profile is the natural partition, separating the desired profile from all other profiles, and paying the minimal needed amount.

Observation 2. *Optimal-cost implementation in Nash equilibrium has partition complexity* 2*. Therefore, feasible implementation in Nash equilibrium also has partition complexity* 2*.*

Proof. Given any game and a profile $s^* \in S$, we show that it is possible to implement s^* as a Nash equilibrium with a contract of partition complexity 2. Moreover, for any $\epsilon > 0$ there is such a contract with ϵ-optimal cost.

Consider the contract that pays 0 to each agent if the profile played is not s^*, and pays every agent i the amount of $t_i = \hat{t}_i^{NE} + \epsilon/n$ if the profile played is s^*. The profile s^* is a Nash equilibrium as given that all agents other then i are playing s^*_{-i}, agent i gets utility $U_i(s_i, s^*_{-i})$ from playing $s_i \neq s^*_i$, while he gets utility $U_i(s^*)+t_i = \max_{\hat{s}_i \in S_i} U_i(\hat{s}_i, s^*_{-i})+\epsilon/n \geq U_i(s_i, s^*_{-i})+\epsilon/n$ from playing s^*_i. The claim that this contract has ϵ-optimal cost is trivial. □

3.2 Dominant Strategies

We have seen that optimal-cost implementation in Nash equilibrium has partition complexity 2, the complexity that is also necessary for feasible implementation. Implementation in dominate strategies is much more demanding and requires linear complexity. Interestingly, for dominant strategies, optimal-cost implementation requires strictly higher complexity than just feasible implementation, although the gap is not large. While for any game it is possible to implement any profile in dominant strategies with a contract of complexity $n + 1$ by just counting the number of agents playing the desired strategy[6], for the effort game such complexity is not sufficed for optimal implementation. In particular, any symmetric contract that just pays the agent according to the number of agents playing the desired strategy cannot achieve optimal-cost implementation (even for a symmetric effort game). Yet, optimal-cost implementation is possible by a refinement of the counting contract that isolates every profile in which exactly one agent does not play the desired strategy. Such a contract has partition complexity $2n$ and is clearly more complex than the coarser contract that only achieves a feasible implementation by counting. The proof actually presents a contact with only $2n - 1$ parts (created by joining the desired profile and all profiles in which none plays the desired strategy).

We remark that feasible implementation in dominant strategies can be achieved with a contract that only uses a *bound* on the payoffs of the game, there is no real need to know the exact details of the game (the space of strategies and the exact payoffs). For optimal implementation Nash payments are also needed, but not more than that.

[6]Note that this is a refinement of the simple contract that implements the profile as a Nash equilibrium, thus the contracts can also be compared in terms of refinement.

Theorem 3. *Feasible implementation in dominant strategies has partition complexity of* $n+1$*. Optimal-cost implementation in dominant strategies has partition complexity of at least* $n+2$ *and at most* $2n-1$*.*

Proof. The claim that feasible implementation in dominant strategies has partition complexity of at least $n+1$ follows from a lower bound proven for a weaker solution concept in Lemma 9. That lemma shows that feasible implementation of a profile as the unique profile that survives iterative removal of strictly dominated strategies has partition complexity of at least $n+1$. To complete the proof of the theorem we present three claims, the theorem follows directly from these claims.

We first present a contract with complexity $n+1$ that implements the desired profile as a profile of strict dominant strategies.[7] Although we present an explicit contract with some specific payments, the contract really only needs to have a bound on the payoffs of the game.

Lemma 4. *Feasible implementation in dominant strategies has partition complexity of at most* $n+1$*.*

Proof. The claim that feasible implementation in dominant strategies has partition complexity of at most $n+1$ follows from the following simple contract with $n+1$ parts. Fix some profile s^* that should be implemented. Each part P_j for $j \in \{0, 1, 2, \ldots, n\}$ includes all strategy profiles in which exactly j players follow what they need to play according to the strategy s^*. Let $Z = 1 + 3 \times \max_i \ \max_{s,s'} (U_i(s) - U_i(s'))$.[8] If the strategy belongs to P_j then all players are paid $j \times Z$. It is easy to verify that every player i has a strict dominant strategy to play s_i^*, as Z is huge and he loses Z by deviating. □

We next show that a contract with complexity $2n-1$ can implement the desired profile as a profile of strict dominant strategies with optimal cost. Now, the contract needs to know both the Nash payments and a bound on the payoffs of the game, but does not really need any more details about the game.

Lemma 5. *Optimal-cost implementation in dominant strategies has partition complexity of at most* $2n-1$*.*

Proof. The claim that optimal-cost implementation in dominant strategies has partition complexity of at most $2n-1$ follows from the following contract with $2n-1$ parts, which is a slight modification of the contract presented in Lemma 4. Fix some profile s^* that should be implemented. Each part P_j for $j \in \{1, 2, \ldots, n-2\}$ includes all strategy profiles in which exactly j players follow what they need to play according to the strategy s^*. Part P_n includes the profiles in which no player is following his desired strategy, and the profile in which all players follow their desired strategy. There are also n additional parts (instead of p_{n-1}), which will be denoted P_{n-1}^i for $i \in [n]$, each part P_{n-1}^i includes all profiles in which all players but i are following the desired strategy, while i does not. If the strategy belongs to P_j for $j \in [n-2]$ then all players are paid $j \times Z$. For part P_n all players are paid slightly above their Nash payment, player i is paid $t_i = \hat{t}_i^{NE} + \epsilon/n$. For P_{n-1}^i for $i \in [n]$, player i is paid 0,

[7] Which by definition makes the profile the *unique* profile of dominant strategies.

[8] A naive computation of Z requires exponential time, but we do not really care about the exact value of Z, all is really needed is that Z is large compare to the gains by any deviation. If payoffs are bounded within a interval, knowing that interval is enough to compute some Z that is good enough for the contract to work.

while every other player $j \neq i$ is paid $t_j = (n-1) \times Z$. It is again easy to verify that every player i has a strict dominant strategy to play s_i^*. □

Finally, we show that optimal-cost implementation in dominant strategies has partition complexity of at least $n+2$, thus require strictly higher complexity than feasible implementation.

Lemma 6. *For the effort game, optimal-cost implementation of the all-effort profile in dominant strategies has partition complexity of at least $n+2$.*

Proof. To prove the claim we show that for the effort game, for $\epsilon > 0$ that is small enough, any contract with partition complexity $n+1$ does not have ϵ-optimal cost, thus any contract that has ϵ-optimal cost has complexity larger than $n+1$. By Lemma 5 there is an ϵ-optimal contract implementing the all-effort in dominant strategies with cost $\sum_i c_i + \epsilon$, as c_i, the cost of effort of i, is also the Nash payment for i.

Fix any contract with $n+1$ parts and assume that it is ϵ-optimal. We can assume without loss of generality that the all-effort profile belongs to P_0. As this is an ϵ-optimal cost, each player i is paid between c_i and $c_i + \epsilon$. Now consider any profile $s^{(i)} = (0, 1_{-i})$, the profile in which all players but i exert effort, while i shirks. Clearly any such profile does not belong to P_0, as i has a dominant strategy to exert effort. We also claim that for $i \neq j$, $s^{(i)}$ and $s^{(j)}$ must belong to different parts. This is so as in the part that includes $s^{(i)}$ it must be the case that the payment to i is at most ϵ (as otherwise i would prefer not to exert effort when all other do), while in the part that includes $s^{(j)}$ it must be the case that the payment to i is at least c_i (otherwise i ends up with negative utility for some profile in which he exerts effort, contradicting exerting effort being a dominant strategy for i). This implies that there are n parts on top of the parts that includes the all-effort profile. Finally, we claim that any profile in which exactly two players shirk cannot belong to any of the above $n+1$ parts. Consider the profile $s^{(i,j)} = (0, 0, 1_{-\{i,j\}})$ for $i \neq j$. We first claim that $s^{(i,j)}$ cannot be in the same part as $s^{(i)}$ as in this case j's does not have a dominant strategy to exert effort. Similarly, $s^{(i,j)}$ cannot be in the same part as $s^{(j)}$. Finally, we consider the possibility that $s^{(i,j)}$ is in the same part as $s^{(k)}$ for $k \notin \{i,j\}$. In that case, the payment for k in this part must be at least c_k, as k exerts effort in $s^{(i,j)}$, but on the other hand, it cannot be more than ϵ, since in the all effort profile k is paid at most $c_k + \epsilon$ and k should have the incentive to exert effort even if all others exert effort. □

This completes the proof of the theorem. □

Observe that the complexity of implementing the all-effort profile in dominant strategies is only $2n-1$ and not 2^n which is the complexity of the naive contact that promises to pay the cost of effort for every agent exerting effort and zero otherwise. Also note that if the contract was restricted to pay at most the cost of effort to an agent, independent of the profile (could not make promises for larger payments when profiles other than the all-effort profile are played), then exponential complexity of 2^n would have been necessary for feasible implementation. Thus, our assumption that off-equilibrium promises are allowed to be large is necessary to achieve contracts with a reasonable partition complexity.

3.3 Iterative Solvability

While optimal-cost implementation in dominant strategies has complexity of at least $n+2$, which is higher than the complexity required for (non-optimal) implantation, once we move to the weaker solution concept of unique profile that survives iterative removal of strictly dominated strategies, it is possible to get optimal-cost implementation with complexity $n+1$, the required complexity even for feasible implementation. Observe that the gap in complexity between this notion and Nash implementation is large (linear vs. constant), while the gap between this notion and implementation in dominant strategies is very small (both linear).

Theorem 7. *Optimal-cost implementation as well as feasible implementation of a profile as the unique profile that survives iterative removal of strictly dominated strategies has partition complexity of* $n+1$.

The theorem follows directly from Lemma 8 and Lemma 9 below.

Lemma 8. *Optimal-cost implementation of a profile as the unique profile that survives iterative removal of strictly dominated strategies has partition complexity of at most* $n+1$.

Proof. Given any game and profile $s^* \in S$, we show that there exists a contract with partition complexity $n+1$ that implements s^* as the unique profile that survives iterative removal of strictly dominated strategies. Moreover, we show that for any $\epsilon > 0$ there is such a contract with ϵ-optimal cost.

We first define the contract. Fix some permutation over the agents. Define P_j for $j \in [n+1]$ to be the set that include all profiles s satisfying $s_i = s^*_i$ for $i < j$ and $s_j \neq s^*_j$. Note that P_{n+1} includes the profile s^*, and only that profile.

Next we define the payments for each part. We first define the payments for part P_j such that $j \in [n]$. The payment to every agent $i \geq j$ at P_j is 0. The payment to every agent $i < j$ at P_j is

$$t_i = \max_{\hat{s} \in S}(U_i(s^*_1, s^*_2, \ldots, s^*_{i-1}, \hat{s}_i, \hat{s}_{i+1}, \ldots, \hat{s}_n) - U_i(s^*_1, s^*_2, \ldots, s^*_{i-1}, s^*_i, \hat{s}_{i+1}, \ldots, \hat{s}_n)) + \epsilon/n$$

Finally, the payment to every agent i at P_{n+1} is

$$t^{NE}_i = \hat{t}^{NE}_i + \epsilon/n = \max_{\hat{s}_i \in S_i}(U_i(\hat{s}_i, s^*_{-i}) - U_i(s^*)) + \epsilon/n$$

We next show that s^* is the unique profile that survives iterative removal of strictly dominated strategies. First observe the agent 1 has a strictly dominant strategy to play s^*_1. Fix any strategies of the other agents $s_{-1} \in S_{-1}$. If $s_{-1} = s^*_{-1}$ then any strategy $s_1 \neq s^*_1$ will result with agent 1 having utility $U_1(s_1, s^*_{-1})$ (as he is paid 0 at S_1) while agent 1's utility when playing s^*_1 is $U_1(s^*_1, s^*_{-1}) + t^{NE}_1 = \max_{\hat{s}_1 \in S_1} U_1(\hat{s}_1, s^*_{-1}) + \epsilon/n \geq U_1(s_1, s^*_{-1}) + \epsilon/n$.

If $s_{-1} \neq s^*_{-1}$ then any strategy $s_1 \neq s^*_1$ will result with agent 1 having utility $U_1(s_1, s_{-1})$ (as he is paid 0 at S_1) while agent 1's utility when playing s^*_1 is $U_1(s^*_1, s_{-1}) + t_1 = U_1(s^*_1, s_{-1}) + \max_{\hat{s} \in S}(U_1(\hat{s}) - U_i(s^*_1, \hat{s}_{-1})) + \epsilon/n \geq U_1(s_1, s_{-1}) + \epsilon/n$.

Next we assume that each agent j from 1 to $i-1$ is playing s^*_j, and need to show that agent i now has a strictly dominant strategy to play s^*_i. The exact same arguments as above hold for the game restricted to each agent j from 1 to $i-1$ is playing s^*_j.

Finally we show that for any $\epsilon > 0$ there is such a contract with ϵ-optimal cost. Indeed the cost of the above contract is $\sum_{i \in N} t^{NE}_i$, while any other contract the implements s^* with the weaker notion of Nash equilibrium has cost at least $\sum_{i \in N} \hat{t}^{NE}_i = \left(\sum_{i \in N} t^{NE}_i\right) - \epsilon$. □

Lemma 9. *Feasible implementation of a profile as a profile that survives iterative removal of strictly dominated strategies has partition complexity of at least $n+1$.*

Proof. To prove the claim we show that for the effort game, any contract implementing the all-effort profile in iterative elimination of strictly dominated strategies has partition complexity of at least $n+1$.

Consider a contract with partition complexity $k+1$, such a contract partitions the strategy profiles to $k+1$ disjoint sets. We show that if there are less than $n+1$ such sets, the contract does not implement the profile as promised.

Let $P_1, P_2, \ldots, P_{k+1}$ be the partition of the space of strategies profiles to $k+1$ sets. (w.l.o.g. all of them are non-empty). Assume w.l.o.g. that agents are sorted by the order of elimination, that is, agent 1 has a dominant strategy to exert effort. Given that agent 1 exerts effort, agent 2 has a dominant strategy to exert effort, etc. We denote the payment to agent i is sets P_j by $p_{i,j}$. Assume without loss of generality that the sets are sorted in some order that satisfies the following property: for any $j > i$ it holds that $p_{i,j} \geq p_{i,i}$. for example, set P_1 is some set for which agent 1 gets the minimal possible payment over all payments that agent 1 gets. After fixing this set, P_2 is some set in which agent 2 get at least as much as he gets in any of $P_3, P_4, \ldots, P_{k+1}$, etc.

We claim by induction that for $i \in \{1, 2, \ldots, k\}$, all profiles in which every agent $j \leq i$ exerts effort, do not belong to any set P_r for $r \leq i$. This implies that the all-effort profile does not belong to P_i is $i \leq n$, and thus there must be at least $n+1$ parts in the partition.

We first prove the claim for $i = 1$. Assume that there is a profile $s \in P_1$ in which agent 1 exerts effort. Given that the other agents are playing s_{-1}, agent 1 gets minimal payment when exerting effort (playing s_1). But effort is costly, and if agent i shirks, he will be paid at least as much but will not bear the cost of effort, increasing his utility. This contradicts the fact that agent 1 has a dominant strategy to exert effort. This proves the base of the induction.

Next we prove the induction step. Assume that the claim holds for every $j < i$, we next prove the claim for i. Consider the process of eliminating dominated strategies, and consider the i-th step in which shirking was eliminated for every agent $j < i$, and now shirking needs to be eliminated for agent i. We know by the induction hypothesis that every profile s in which every agent $j < i$ exerts effort belongs to some P_r for $r \geq i$. The contract must ensure that i is better off exerting effort in this case. But effort is costly and i can ensure a payment of at least $p_{i,i}$ even when shirking (as for any $r > i$ it holds that $p_{i,r} \geq p_{i,i}$), thus for agent i to have a dominant strategy to exert effort (given that every $j < i$ exerts effort), any profile in which every agent $j \leq i$ exerts effort does not belong to P_i. □

3.4 Unique Nash Equilibrium: Pure and Mixed

In this section we consider implementing a profile as the *unique pure Nash equilibrium*, or as the *unique Nash equilibrium* (among both pure and mixed).

3.4.1 Unique Pure Nash

While feasible implementation of a profile as a Nash equilibrium has complexity of only 2, such complexity is not enough to ensure that the profile is the *unique* Nash equilibrium profile. We show that complexity of 3 is required to achieve this, and is also sufficient for optimal-cost implementation.

Theorem 10. *Feasible implementation of a profile as the unique pure Nash equilibrium has partition complexity of at least* 3*. Optimal-cost (and thus also feasible) implementation of a profile as the unique pure Nash equilibrium has partition complexity of* 3*.*

The theorem follows from the following two claims.

Lemma 11. *For the effort game, any contract with partition complexity* 2 *cannot implements the all-effort profile as the unique pure Nash equilibrium.*

Proof. Assume that such an implantation is possible, we will derive a contradiction. Assume that we are given a partition to only two parts, call them P_1 and P_2. Without loss of generality assume that all-effort vector belong to P_1. Let p_i denote the payment to agent i if the vector is in set P_1, and let q_i denote the payment to agent i if the vector is in set P_2. To get the all-effort vector to be a Nash equilibrium, in P_1 each agent i must be paid p_i that is at least c_i. If a vector with agent i shirking and all other $n-1$ agents exerting effort belong to P_1 then the all effort is not an equilibrium: i would rather shirk as this will increase his utility because he will have less cost but will receive the same payment of p_i. We conclude that all vectors with exactly $n-1$ agents exerting effort belong to P_2. The payment q_i to agent i in P_2 must be at most $p_i - c_i$, as otherwise that agent shirking from the all-effort vector would be a beneficial deviation, thus $0 \leq q_i \leq p_i - c_i < p_i$.

Now consider the all-shirking vector. It must be in P_2 as if it belongs to P_1 it is an equilibrium: each agent i gets utility of $p_i \geq c_i > 0$ and no deviation can increase that (as for both P_1 and P_2 the payment to i does not exceed p_i, as $q_i < p_i$, and exerting effort has higher cost than shirking). As the all-shirking vector is not an equilibrium, a beneficial deviation in which one agent exerts effort must exist. Assume that agent is i, it must be the case that the vector when i exerts effort and all other shirk belongs to P_1 (since deviating from shirking to exerting effort cannot be beneficial if the payment does not change). For this vector not to be an equilibrium, another agent $j \neq i$ must have a beneficial deviation of exerting effort, but this new vector cannot belong to P_2 (as in P_2 agent j's payment q_i is smaller than p_i and deviating by shirking increases the cost), and it cannot belong to P_1 as the payment is the same and the cost increases. Thus we have a contradiction and this concludes the proof. □

We next show that optimal-cost implementation of a profile as the unique pure Nash equilibrium has partition complexity of 3. The contract we built essentially put players 1 and 2 into a matching pennies game which has no pure Nash equilibrium. This game is slightly modified to make sure the desired profile is indeed an equilibrium.

Lemma 12. *Optimal-cost implementation of a profile as the unique pure Nash equilibrium has partition complexity of* 3*.*

Proof. Fix any game and let s^* be the profile that we would like to implement as the unique pure Nash equilibrium. Let $Z = 1 + 3 \times \max_i \max_{s,s'} (U_i(s) - U_i(s'))$. Consider the contract with three parts P_0, P_1, P_2 as follows. Part P_0 includes only the profile s^*, each player i is paid slightly above his Nash payment, player i is paid $\hat{t}_i^{NE} + \epsilon/n$. Part P_2 includes any profile s such that $s_1 \neq s_1^*$ while $s_2 = s_2^*$, and any profile s such that $s_1 = s_1^*$ while $s_2 \neq s_2^*$, excluding the profiles in which every player $i \neq 2$ plays s_i^*. For profiles of part P_2, player 2 is paid Z, while all other are paid 0. Part P_1 include every other profile. For profiles of part P_1, player 1 is paid Z, while all other are paid 0.

We claim that s^* is the unique pure Nash equilibrium in the game induced by this contract. To prove this we first observe that s^* is indeed a Nash equilibrium as any player i deviating to s_i will end up with 0 payment and utility $U_i(s_i, s_{-i}^*)$, while playing s_i^* gives i utility of $U_i(s^*) + \hat{t}_i^{NE} + \epsilon/n = \max_{\hat{s}_i \in S_i} U_i(\hat{s}_i, s_{-i}^*) + \epsilon/n \geq U_i(s_i, s_{-i}^*) + \epsilon/n > U_i(s_i, s_{-i}^*)$. The claim that this contract has ϵ-optimal cost is trivial, as paying the Nash cost is necessary for Nash implementation (even non-unique).

We then show that for any profile that is not obtained by a single player deviating from s^*, if the profile is in P_1 then player 2 has a beneficial deviation, and if the profile is in P_2 then player 1 has a beneficial deviation. For such $i \in \{1, 2\}$, if the profile s is such that i is not playing s_i^*, the deviation is to switch to s_i^*, and if i is playing s_i^*, the deviation is to switch to a strategy s'_i which maximizes $U_i(s'_i, s_{-i})$ over all $s'_i \neq s_i^*$. In any case the deviation will result with player i being paid Z. Z is picked to be so big such that no matter what the others are playing, a player would always prefer being paid Z than playing any strategy and being paid 0. □

3.4.2 Unique (Mixed) Nash

A unique profile that survives iterative removal of strictly dominated strategies must be the unique Nash equilibrium of a game. Additionally, it is clear that the cost of implementing a profile as the unique Nash equilibrium is at least the Nash cost. Thus, the contract of Lemma 8, which has Nash cost and implements a profile as the unique one that survives iterative removal of strictly dominated strategies, provides an upper bound on the complexity of optimal implementation as the unique Nash equilibrium.

Corollary 13. *Optimal-cost implementation of a profile as the unique mixed NE has partition complexity of at most* $n + 1$*.*

The exact complexity of both feasible and optimal-cost implementation of a profile as the unique mixed NE is left as an open problem.

3.5 Implementation in Undominated Strategies

Recall that for the effort game, it is a dominant strategy for the players to shirk, thus implementing this profile in undominated strategies must have partition complexity of at least 2. We next show that any profile can be implemented with 0 cost by a contract of partition complexity 2, which is the best the principal can hope for. Such cost is lower than the cost of implementing a profile as a Nash equilibrium. The proof idea is to make the desired strategy of each player undominated without ever actually paying any agent when they all play the desired profile. To achieve this the contract promises high payment for the desired strategy of each player but only when other players are deviating. When none deviates, no payments are made.

Observation 14. *Optimal-cost (and thus also feasible) implementation in undominated strategies has partition complexity* 2*.*

Proof. Given any game and a profile $s^* \in S$, we show that it is possible to implement s^* in undominated strategies with a contract of partition complexity 2. The implementation we construct has zero cost, and thus is clearly optimal.

Let $Z = 1 + 3 \times \max_i \ \max_{s,s'} (U_i(s) - U_i(s'))$. The partition has two parts. In the first part we have all profile in which exactly one player is following the desired strategy. For this partition every player is paid Z. The other part includes all remaining profiles (including the profile s^*), and the payment to any player is 0. Observe that in the case that all players but i are not following the desired profile, i's unique best response is to follow his desired strategy s_i^*, thus this strategy is undominated, and the contract indeed implements s^* in undominated strategies. □

4 Coalitional Deviation Concepts

In this section we consider coalitional deviation concepts: strong Nash equilibrium and unique strong Nash equilibrium.

4.1 Strong Nash Equilibrium

As a strong Nash equilibrium is a stronger solution concept than Nash equilibrium, clearly the partition complexity of such an implementation is at least 2. It turns out that although this is a much stronger solution concept, the complexity of an optimal-cost implantation of strong NE is still 2. The proof idea is to show that there is an optimal contract that never pays any agent unless the desired profile is played. This is so as if we take any optimal contract that implements some desired profile as a strong Nash equilibrium, and zero all payments when an undesired profile is played, the new contract also implements the desired profile as a strong Nash equilibrium, and has the same cost. We note that this cost might be higher than the Nash cost.

Theorem 15. *Optimal-cost (and thus also feasible) implementation of a profile as a strong Nash equilibrium has partition complexity* 2*.*

Proof. We show that given any game and profile $s^* \in S$, it is possible to implement s^* as a strong Nash equilibrium with a partition with only 2 parts, an moreover, for any $\epsilon > 0$ there is such a contract with ϵ-optimal cost.

Consider contracts that can specify payment $C_i(s)$ for any agent i and strategy profile $s \in S$. We claim that there is such a contract that pays every agent 0, unless s^* is played. This clearly implies that the partition complexity is 2, as one part in this contract includes s^* and the other includes all other profiles.

Fix any contract (with arbitrary complexity) C^* that implements s^* as a strong Nash equilibrium and has ϵ-optimal cost (such a contract exists by definition). Now consider modifying this contract that pays 0 unless s^* is played (and the same payments as C^* otherwise). As all payments are non-negative, removing these payments only decreases the incentive of coalitions to deviate from s^*, thus s^* remains a strong Nash equilibrium with the modified contract. □

4.2 Unique Strong Nash Equilibrium

While the complexity of feasible implementation of a profile as the *unique* Nash equilibrium is higher than the complexity feasible implementation of a profile as a Nash equilibrium, this is not the case for strong Nash equilibrium, as we show next. This is so as increasing payment for the desired profile does not help in uniqueness of Nash equilibrium, but does so for strong Nash equilibrium. If payments for the desired profile are high enough, no other profile can be a strong Nash equilibrium as all agents would rather deviate to the desired profile together, while for Nash equilibrium such high payments do not rule out the existence of other equilibria as only individual deviations are considered. Yet, such a low complexity contract seems to have higher than necessary cost. We conjecture that optimal-cost contracts have strictly higher complexity, and leave the complexity of optimal cost implementation of unique strong Nash equilibrium as an open problem.

Theorem 16. *Feasible implementation of a profile as a unique strong Nash equilibrium has partition complexity* 2*.*

Proof. We show that given any game and profile $s^* \in S$, it is possible to implement s^* as the unique strong Nash equilibrium with a partition with only 2 parts, one only includes s^* with very high payments, and the other include all other profiles, with no payments.

Formally, if s^* is played then each agent i is paid $Z_i = 1 + 3 \times \max_{s,s'} |U_i(s) - U_i(s')|$. It is easy to verify that this is indeed a strong Nash equilibrium (as the loss of payment of Z_i for every deviating player i eclipses any gain from the game) and that any other profile is not a strong Nash equilibrium as it is beneficial for all agents not playing according to s^* to deviate together such that the final profile will be s^*. □

We note that the contract presented in Theorem 16 also creates a feasible implementation of a profile as a unique Coalition Proof Nash Equilibrium.

5 Conclusions

In this paper we have studied the complexity required for the implementation of multi-agent contracts under a variety of solution concepts. Our focus was on the complexity that arises from verifying the contract, complexity which increases as the contract is refined. As not every two contracts are comparable with respect to refinements, we extend the partial order to a complete order by considering the size of the output set as our complexity measure. We found that there is a large gap between the complexity of weak solution concepts like Nash equilibria and strong concepts like dominant strategies. Additionally, we found that the complexity of feasible implementation is usually also sufficient for optimal implementation, event when insisting on uniqueness.

Clearly, our complexity measure is only one of many possible measures for complexity of contracts. One issue that it does not captured by our complexity measure is the length of the contract's description (say when written in a "reasonably simple" language). Trying to address such an issue, one might specify a small set of queries that the contract is allowed to use in its contingencies (e.g., queries like: "did agent i play s_i^*?") and define the complexity to be the minimal number of such queries in the contract, or in the longest chain of queries that need to be asked for a arbitrary profile. An interesting direction for future work is to study such measures and compare between the results achieved for different measures.

Our analysis is confined to complete information implementation. In many settings, agents have private types and the game is actually a Bayesian game. For example, the cost of effort in an effort game might be private information. A natural direction for future research is to consider the complexity of implementation in the Bayesian setting (incomplete information). While for the full information setting the actions space and the strategy space are the same, this is no longer the case in the Bayesian setting, as now strategies are mappings from types to actions. It is natural to assume that the principal can only condition payments upon actions (and not the unobserved strategies), yet he would like to implement some strategy profile in a given solution concept. For example, in a single-item auction setting, the principal might like agents to have a dominant strategy to be truthful about their value for the item. While in full information such an implementation was always possible, this is no longer the case in the Bayesian setting, but when it is, it is interesting to understand the complexity of feasible and of optimal implementations.

Acknowledgments The authors are grateful to Noam Nisan for very helpful discussions.

References

[1] Anderlini, L. and Felli, L. 1994. Incomplete written contracts: Undescribable states of nature. *The Quarterly Journal of Economics 109,* 4, 1085–1124.

[2] Anderlini, L. and Felli, L. 1999. Incomplete contracts and complexity costs. *Theory and Decision 46,* 1, 23–50.

[3] Babaioff, M., Feldman, M., Nisan, N., and Winter, E. 2012. Combinatorial agency. *Journal of Economic Theory (JET) 147,* 3, 999–1034.

[4] Dye, R. A. 1985. Costly contract contingencies. *International Economic Review 26,* 1, 233–50.

[5] Hart, O. D. and Moore, J. 1988. Incomplete contracts and renegotiation. *Econometrica 56,* 4, 755–85.

[6] Hart, S. and Nisan, N. 2013. The menu-size complexity of auctions. In *ACM Conference on Electronic Commerce.* 565–566.

[7] Holmstrom, B. 1982. Moral hazard in teams. *Bell Journal of Economics 13,* 2, 324–340.

[8] Holmstrom, B. and Milgrom, P. 1987. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica 55,* 2, pp. 303–328.

[9] Kolmogorov, A. N. 1998. On tables of random numbers. *Theor. Comput. Sci. 207,* 2, 387–395.

[10] Li, X. and Yao, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences 110,* 28, 11232–11237.

[11] Melumad, N., Mookherjee, D., and Reichelstein, S. 1997. Contract complexity, incentives, and the value of delegation. *Journal of Economics and Management Strategy 6,* 1, 257–289.

[12] Segal, I. 1999. Complexity and renegotiation: A foundation for incomplete contracts. *Review of Economic Studies 66,* 1, 57–82.

[13] Wang, Z. and Tang, P. 2013. Optimal mechanisms with simple menus. *CoRR abs/1311.5966*.

[14] Williamson, O. 1975. *Markets and hierarchies, analysis and antitrust implications: a study in the economics of internal organization*. A study in the economics of internal organization. Free Press.

[15] Williamson, O. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press.

[16] Winter, E. 2004. Incentives and discrimination. *American Economic Review 94,* 3, 764–773.