1

Reciprocal Contracting

Michael Peters
University of British Columbia
997-1973 East Mall, Vancouver, BC, V6K1Z1
604-822-4418
peters@econ.ubc.ca

Abstract

This paper models competing mechanism games as extensive games where the extensive form is incompletely understood by a modeler, typically because the modeler doesn't see all the messages that are being exchanged and doesn't understand all the contracts that can be enforced. For this reason, the revelation principle can't be used to characterize supportable outcomes. The paper provides a relatively weak restriction, referred to as regularity, on the unknown part the competing mechanism game. This condition makes it possible to characterize the set of supportable equilibrium outcomes of the unknown game using information about the part of the game the modeler does understand. In addition, the paper provides a canonical game called the reciprocal contracting game which supports as an equilibrium every equilibrium outcome of any regular competing mechanism game that embeds the known part of the game. As a consequence, the reciprocal contracting game can be used as a stand-in for the true game.

1. Introduction

One of the pricing techniques that is fairly common on the internet is something called *click stream pricing.* For example, a recent request for a price quote for a generic desktop computer offered to bundle the computer with a two year warrant. Below the price quote and its associated "Add to Cart" button was a check box with an offer to buy a 2 year warranty for $79. Checking the box had the effect of adding the warranty to the cart. Leaving the box unchecked then clicking the checkout button, brought up a page with the usual total charges. Below this total charge was a banner saying "People who purchased this Item also liked the Following items:". There again was the 2 year warranty, but now at a price of $56.

[1] I would like to thank the editor, three referees and numerous seminar participants for their helpful comments. This paper was completed with support of the Social Sciences and Humanities Research Council of Canada.

This example illustrates a number of things about the way modern digital markets operate (and how different they are from competitive markets as envisaged by Arrow-Debreu). First, the transaction itself doesn't involve a seller - just a dumb computer program that comes up with a price offer. This program can't negotiate or change its mind. However, what it does depends on messages it receives from the buyer. In the click stream pricing case, the message is a stream of urls that the buyer visited before asking for a quote. The actual offer would depend on cookies that had been placed on the buyer's browser, whether the buyer had 'logged in' to his or her account with the seller, and so on. The quote might even depend on whether the buyer had visited other websites.[2] This kind of price discrimination will occur even when there is intense competition between suppliers and goods produce only private values (like warranties).

What makes these markets hard to understand theoretically is that many of the details associated with the mechanisms that sellers use are buried in complex computer networks. It is almost impossible to describe a digital market by listing the mechanisms that players use along with the messages that these mechanisms entail. Moreover, the process by which sellers process messages is continually changing since there is lots of innovation in information processing on the internet. Whatever mechanisms sellers are allowed to use in a model are likely to be obsolete by the time the model produces a theorem.

Digital markets then provide a setting where it is necessary to understand what is feasible in equilibrium without knowing exactly the extensive game traders are playing. A broader approach is needed. The purpose of this paper is to try to provide one. We start with a a formal definition of competing mechanism games that is broad enough to encompass any imaginable digital market. We then provide a relatively weak condition on these games, referred to as *regularity,* that makes it possible to characterize the outcome functions that can be supported as perfect Bayesian equilibrium of games that embed the known parts of the game - referred to here as the *default game*. We provide a complete characterization of this set of outcome functions. The significance is that equilibrium of games that are incompletely known can be understood using information about the part of the game that is known.

Finally, we provide a single game, called the reciprocal contracting game, that has the property that an outcome function can be supported as a perfect Bayesian equilibrium in a regular competing mechanism game if and only if it can be supported as an equilibrium for the reciprocal contracting game. The advantage of this is that it isn't necessary to have a whole class of games to understand what is supportable with competing mechanisms in the way that the entire class of direct mechanisms is needed to understand supportable outcomes in mechanism design.[3]

[2]Twitter and Facebook can place cookies on my computer when I load the "Like" or "Follow us on" buttons on other websites into my browser. The dumb program can (presumably for a fee) query Twitter to check whether the same browser has visited a competitors website.

[3][1] for example, shows that the set of incentive compatible and individually rational outcomes in a Bayesian game can be described using a class of commitment games.

The basic logic of equilibrium in the reciprocal contracting game is very simple - though many of the details concerning communication can be complex. A competing mechanism game works in much the same way as a repeated game. Yet, the characterization of potential equilibrium is much more precise in competing mechanism games than in repeated games when there is incomplete information. Indeed, the set of outcome functions that can be supported as *perfect* Bayesian equilibrium in regular competing mechanism games is the set of incentive compatible and individually rational outcome functions as they are defined in Myerson's textbook ([2]).

The reciprocal contracting game itself is a regular competing mechanism game. In particular what that means is that players commit their own actions based on messages they exchange with other players. In this sense, the players are using mechanisms that are conceptually no more complicated or demanding than auctions.

The fact that there is a decentralized game like the reciprocal contracting game that supports *all* incentive compatible and individually rational outcomes at the same time is troubling.

What it means is that very good default games can support very bad equilibrium outcomes when the parts of the game the modeler doesn't understand allow traders to circumvent the rules of the default game. As argued above, this seems to be the case in digital markets. The fact a Myerson like coordinator can design a particular direct mechanism that implements some desirable incentive compatible and individually rational outcome does not mean that players can't undo the outcome by manipulating it using their own contracts. So even in environments where there does exist a coordinator who can give binding instructions to agents about how to behave, there is a question about what outcome functions one should expect.

Furthermore, a Myerson like coordinated outcome seems implausible in modern markets for two reasons. First, traders in digital markets have very well defined property rights. Except under special circumstances, it is a seller's right to charge whatever price he or she likes for the good or service they sell. Of course, to write a computer program to generate a price offer, sellers have to be able to make whatever commitments they want with respect to how their price offers depend on messages. On the other side of the market, a buyer cannot be coerced into a trade if they think the price is too high. For a Myerson like coordinator to be effective, both these rights have to be taken away - buyers and sellers must typically agree to a trade before they know the terms.

Furthermore, the set of allocation rules that are interim incentive compatible typically depend on beliefs, which may be constantly changing in some markets. A Myerson coordinated market would then require constant rejigging of the rules of trade to make the market effective. This doesn't seem to occur - typically the rules of trade change very slowly over time. Perhaps another way to put this is that the rules and trading conventions that we take for granted in markets are precisely those that haven't change for a long time.

The reciprocal contracting game preserves all these rights because it is *decentralized* in exactly the sense that it takes certain property rights for granted (those defined by the default game) but allows players to make commitments about *their own* actions based on messages. It is perhaps surprising that what appear be constraints on the ability of the players to coordinate their actions through a centralized coordinator has no implication for the set of outcome functions that can be supported as perfect Bayesian equilibrium.

[3] and [1] point out that the Myerson coordinator can be interpreted in a different way. Once the mapping from types to actions is known - for example, an incentive compatible and individually rational mapping, one can imagine finding the direct mechanism that implements this outcome function, then re-declaring the types to be commitment devices. The revelation principle can then be understood as defining a game in commitment devices, or what Forges calls a *commitment game.* It follows immediately from this that if an outcome function is incentive compatible and individually rational in Myerson's sense, there is *some* commitment game which can replace the default game so that the desired outcome is a Bayesian Nash equilibrium of this new default game. Since this is essentially a reinterpretation of the revelation principle, the comments made about the revelation principle above apply to this approach as well.

A Short Digression on the Literature

One of the things that makes modern digital markets so different from classical markets is that is is easy for sellers to interpret messages from buyers. If these messages change when market conditions do, then sellers can use them to respond to the actions of their competitors, even when a market is large and otherwise competitive looking. This is an old idea (for example [4]). It is an important idea because it undermines the the notion of type that is so important in mechanism design. If a buyer's type report is supposed to convey information about what other sellers are doing, and the other sellers' actions depend on the mechanism being used by the first seller, then buyers' types would appear to depend on the mechanism the seller is using. This takes away the revelation principle as a means of comparing the performance of different mechanisms.

The primary contribution of [4] was to show how to define type in such a way that it did not depend on the mechanism that is offered. The method involved creating messages that would illustrate what competitors would do under all kinds of different circumstances. It isn't too hard to see that these messages are necessarily complex. Though the result showed how one could still use the revelation principle to analyze mechanisms in a competitive environment, it also showed that direct mechanisms would be too complex to be of any practical use, either theoretically or in applications.

One approach to this conundrum was provided by the literature on common agency ([5, 6, 7], [8], [9], [10]). In particular, it suggested abandoning the revelation principle and modeling competition by having mechanism designers offer their agents menus instead of direct mechanisms. This approach had two big advantages apart from the fact that menus are conceptually simple. First, it showed why competing mechanisms would typically lead to multiple equilibrium outcomes ([11]), and second, it gave a method that could be used to easily characterize the very attractive outcome functions that could be supported under competition ([12] or [13]).

It still presented two challenges. First, common agency assumes a single agent. With multiple agents a menu offered to a single agent often doesn't make sense.[4] Second, though it is clear in common agency models that there are multiple equilibria, it isn't clear what the set of equilibrium outcomes looks like. This made it hard to assess equilibrium selection arguments.

The extension of the menu approach to multiple agents was suggested by [14]. His idea was to have principals offer their (multiple) agents a menu consisting of a collection of direct mechanisms feasible for the principal. He called a mechanism where agents were offered a menu of direct mechanisms a recommendation mechanism. If the agents could agree on one of the direct mechanisms in the menu[5], the principal would commit himself to implement it. The idea was very simple. If agents saw that all principals were using recommendation mechanisms, then the agents would all recommend a particular direct mechanism. They would have incentives to do this because they expect all the other agents to make the same recommendation. If some principal were to deviate, then the agents would again all coordinate on a mechanism that punished the deviating principal.

This approach makes it apparent that the set of equilibrium outcomes supportable in a competing mechanism game must actually be huge. The only question that seemed to remain is what this set looked like. The paper by [15] provided the first bit of an answer, while at the same time extending the approach to general games (in which all players in a game acted as both principals and agents. Given any Bayesian normal form game as a default game, it is possible to create a 'competing mechanism' game on top of this normal form game by allowing players to write contracts that make commitments to their actions in the default game. [15] used the approach in [14] to show that the set of outcome functions supportable as Bayesian equilibria in the competing mechanism game is always at least as large as the set of incentive compatible and individually rational outcome functions as defined in Myerson's text ([2]).

Yamashita's logic makes part of this clear. When all the principals offer their agents a full menu of direct mechanisms and just ask for a recommendation, they are implicitly agreeing to participate in a centralized direct mechanism. An agent who deviates is effectively refusing to participate. As a consequence, the agents jointly recommend a punishment. Since the non-deviators are committed to carry this punishment out, they could simply recommend the same punishment that the Myerson coordinator would have recommended in the event that some player refused to participate.

The subtle difference between the result in [15] and Myerson is that when the Myerson coordinator punishes a player who doesn't participate, he or she is

[4]Though, the approach is still useful in very specialized environments.

[5]Agreement means specifically that all but at most one of the agents recommends the same mechanism.

constrained to implement a punishment that is independent of what the non-participant decides to do. In a Yamashita like recommendation game, this isn't the case. The non-deviators can see the mechanism that the deviator has chosen, and could therefore recommend a punishment that is tailored to the crime. The implication is that there will be equilibrium outcomes in a recommendation game in which some players' payoffs will be lower than their individually rational payoffs (in Myerson's sense).[6] A complete characterization of the set of outcomes supportable as Bayesian equilibrium in a general recommendation game is as yet unknown.

This complication implies that a full characterization of the equilibrium set requires a restriction on the set of feasible mechanisms. This problem was addressed in [16]. Rather than dealing with the recommendation game, they allow contracts that can depend directly on the contracts of other players (as we do in this paper). In their approach, a contract is just a finite text describing the player's commitment. A contract that depends on the contracts of the other players is a function that maps from the texts of the other players' contracts into a subset of the action space of the player. As in Yamashita, these functions make it possible in principle for one contract to punish a deviating player in a way that depends on what the deviating player does. What they do is to show that if players are allowed to use *definable* functions, then players who want to deviate will always be able to write their (deviating) contracts in such a way the other players cannot condition their actions on the outcomes that the deviating contract ultimately chooses. This restores the min-max structure of the equilibrium payoffs, which allows them to characterize the full set of supportable outcome functions.

There are a number of restrictions in [16] that limit the generality of their result. The environment has to be finite, the characterization is only for pure strategy equilibrium, and mechanisms cannot be random. Perhaps more important, their contracts do not allow any communication of type information apart from that embodied in the contract offer a player makes. For this reason it is hard to compare the set of equilibria with the set of outcome functions that are individually rational and incentive compatible in Myerson's sense, apart from the fact that in the appropriately restricted environment, the equilibrium set is a subset of the incentive compatible individually rational set.

This paper is intended to add to this in a number of ways. First, it is intended to extend the result in [16] to eliminate the pure strategy restrictions and non-random mechanisms and to allow agents to communicate information to principals within a mechanism as they do in a direct mechanism. This makes it possible to give a quite general characterization result. To do this, we cannot use the definable function approach in their paper, since definable functions require everything to be countable. Nonetheless, a characterization requires

[6] The problem is analogous to the problem of figuring out the difference between min-max and maxmin payoffs. In games of complete information, the relationship is simple and well known. The same cannot be said for games of incomplete information.

a restriction on the set of feasible contracts as described above. Reciprocal contracts are conceptually much simpler than definable functions and make it much easier to see how contracts condition on one another.

Second, it should be clear from the discussion above, that none of the games described so far are intended to be descriptive. The messages players are supposed to convey are too complex to be used in practice. The same criticism could be made of the revelation principle since having players announce types is typically infeasible. The point of the papers above is to work out what is feasible, so that implicit restrictions involved in more descriptive models (or indirect mechanisms) can be understood.

In this sense, the reciprocal contracting game described below is intended to play the role of the revelation principle, but for competing mechanism games instead for simple mechanism design problems. For the standard revelation principle, it is clear exactly what indirect mechanisms can be represented as direct mechanisms. It is harder to do this for competing mechanisms because the analog of an indirect mechanism is a game among mechanism designers, typically an extensive form game.

The second contribution here is then to describe the class of extensive games that can be represented by reciprocal contracting. We refer to these games as 'regular' contracting games. To the extent that regular contracting games seem reasonable, reciprocal contracting has an unexpected advantage. Two different indirect mechanisms have very different representations as direct mechanisms. Indeed, two different equilibrium outcomes of the same indirect mechanism would have different representations as direct mechanisms. With reciprocal contracting, two different regular extensive form games will support different equilibrium outcomes which can both be supported by the *same* reciprocal contracting game. This is what makes it possible to use reciprocal contracting to understand environments like internet markets where the exact nature of the extensive games is unknown.[7]

Finally, in most of the literature on this topic, the solution concept is Bayesian equilibrium.[8] Competing mechanism games are extensive form games, so the appropriate solution concept is perfect Bayesian equilibrium. One might expect that refinements might limit the set of supportable equilibrium outcomes, making characterizations based on the revelation principle invalid. In particular, the revelation principle requires the implementation of a punishment for non-participating players that the participants might not want to carry out.

Though it is certainly true that refinements will typically limit the set of outcomes supportable in a given game, we show that the set of regular competing mechanism games in large enough to compensate for this. In the theorem below the solution concept for both regular competing mechanism games *and* the

[7]Note the difference between this and the approach in [1] or [3] where each extensive form game must be represented by a different game in commitment devices.

[8]The exceptions are [14], where there is no real characterization result, and [16] who provide a very weak refinement of Bayesian equilibrium.

reciprocal contracting game is perfect Bayesian equilibrium. To do this, the reciprocal contracting game exploits two properties of competing mechanism games. The first is just commitment, which is fundamental to mechanism design. The second is something which is unique to competing mechanism games - to understand all the equilibria in competing mechanism games, players must be allowed to use messages that convey no information at all about their payoff types. Though this insight was made many years ago by [4], the reciprocal contracting game converts the complex messages they needed into conceptually simple messages that are added to messages about payoff types.

2. Incomplete Information Games and Mechanism Design

The basic approach in what follows is to add the contracting game on top of a basic default game of incomplete information as in the example above. In default game, there are n players. Each player has a finite action set A_i and a finite type set T_i. In standard notation A, A_{-i} represent cross product spaces representing all players actions and the actions of all the players other than i, respectively. Similarly, define $T = \prod_i T_i$, and $T_{-i} = \prod_{j \neq i} T_j$. Types are jointly distributed on T according to some common prior.

Let q be a mixture over the set of action profiles A, i.e., $q \in \triangle(A)$. For any action profile a, we write q_a to be the probability of a under q, and $q_{a_i} = \sum_{a_{-i}} q_{a_i, a_{-i}}$. We use notation q_{A_i} to represent the marginal distribution over A_i and $q_{A_{-i}}$ to be the marginal distribution over A_{-i}. We assume that players have expected utility preferences over lotteries. Then players preferences are given by $u_i : \triangle(A) \times T \to \mathbb{R}$ where u_i is linear in q. An *outcome function* is a mapping $\omega : T \to \triangle(A)$. So player i's payoff from this outcome function is $\mathbb{E}\{u_i(\omega(t), t) | t_i\}$.

Notice that an outcome function here is a description of a set of actions that players use in the default game rather than a set of outcomes in the usual sense. Of course, the default game itself is an indirect mechanism. Ultimately the point of all this is to evaluate different default games. However, at this stage, this game is taken to be fixed.

An outcome function ω is implementable (by a mechanism designer) if the usual incentive compatibility and individual rationality conditions hold. Formally, an outcome function ω is *incentive compatible* if for every i, t_i and t_i',

$$\mathbb{E}\{u_i(\omega(t), t) | t_i\} \geq \mathbb{E}\{u_i(\omega(t_i', t_{-i}), t) | t_i\}. \quad (1)$$

This is completely standard so there is no need to discuss it further. What is different in our approach is what happens when a player refuses to participate in the mechanism designer's scheme. In that case, instead of an exogenous payoff, it will be assumed that the default game takes precedence.

When some player unilaterally refuses to participate, we allow the mechanism designer to implement a 'punishment' involving a mixture over profiles of actions of the participating players.

Let $\rho^i : T_{-i} \to \triangle(A_{-i})$ be an outcome function that is implemented when player i chooses not to participate in the mechanism that implements ω. We refer to this outcome function as a *punishment.* The outcome function ω is *individually rational* if there is a collection of punishments $\left\{\rho^i\right\}_{i=1,n}$ such that for every player i,

$$\mathbb{E}\left\{u_i\left(\omega\left(t\right),t\right)|t_i\right\} \geq \max_{a_i} \mathbb{E}\left\{u_i\left(a_i, \rho^i\left(t_{i-i}\right),t\right)|t_i\right\}. \tag{2}$$

This is just the constraint used by [2] to describe *Bayesian* equilibrium outcome functions for a specific 'default game' of incomplete information. What we are going to show is that, provided there are enough players, the same set of outcome functions can be supported as *Perfect Bayesian equilibria* in some *regular competing mechanism game* that embeds the default game. Notice a number of things - Bayesian equilibria in the default game is replaced by Perfect Bayesian equilibrium in a regular competing mechanism game. The punishments in Myerson don't have to be sequentially rational, while they do here.

Secondly, the equivalence only holds for *regular* competing mechanism games. The next section describes competing mechanism games more generally, and introduces the concept of regularity. Also, it is possible to implement all the outcome functions supportable in some conceivable competing mechanism game with one special, but somewhat abstract, game. This 'game' is defined after the next section. The formal statement of the results follow. The proofs of the main theorems are in the appendix.

2.1. What is a Competing Mechanism Game

One of the things that plagues the competing mechanism literature are the very complex definitions required to describe specific competing mechanism games when market information is being used by players. The description of universal types in [4] is very abstract, while the games described in [14] or [15] are hard to describe because of the complex communication that occurs between players. The common agency literature (for example[5] or [8] or [11]) describe simpler games by focusing on very special kinds of contracts. All those references involve one shot games in which contracts are offered simultaneously. A useful description of competing mechanism games should also capture models in which mechanisms are offered sequentially, as in [6] or privately as in [17] or [18].

A somewhat simpler approach is to model competing mechanism games as extensive form games of incomplete information. In the familiar formalism (for example, [19]) an extensive game is a directed graph consisting of nodes and edges connecting the nodes. We interpret the nodes of this game as opportunities for players to send messages and the edges as the messages themselves. A path through the game (graph) is an ordered sequence of messages.

The meaning of the messages depends on the specifics of the model. Some messages convey commitments, some type information, while some are just cheap talk. In order to interpret the messages, we use an outcome function λ which assigns a profile of actions to each path through the game tree. The profile of actions indirectly determines each player's payoff.[9]

The picture that follows illustrates the reciprocal contracting version of a simple prisoner's dilemma game augmented by contracts that condition on other contracts (with the cheap talk part left out to make it simpler). The players reveal their commitments by sending commitment messages.

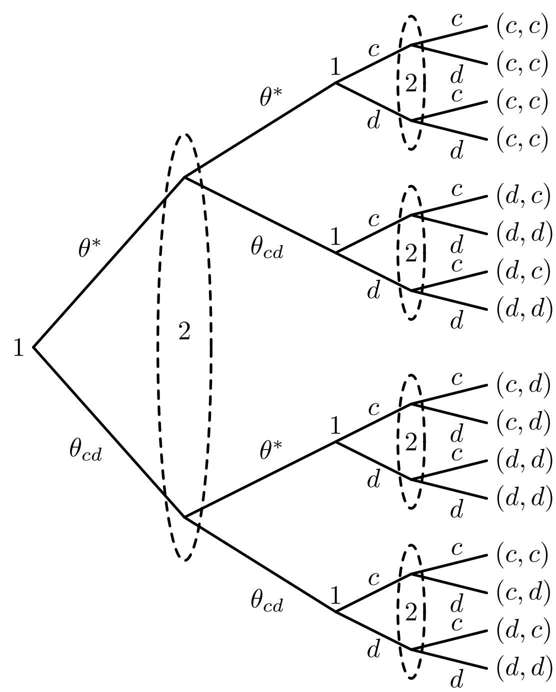

In this game, players announce public messages representing commitments over two rounds. In the first round, each player can offer a contract that conditions directly on the other player's contract. This is the contract θ^* in the picture. The alternative is a contract θ_{cd} that allows the player to defer his choice until the second round. The outcome function λ is displayed on the far right of the picture. Notice that if player 1 announces the message θ^* in his first information set, then the outcome function forces him to use action c in every history in which player 2 uses message θ^*, and to use action d in every other history following that choice. So from the outcome function λ, the interpretation of the message θ^* is that it is a reciprocal contract that commits player 1 to use action c if player to sends signal θ^*, and to use d otherwise.

[9]The fact that messages only determine profiles of actions is what distinguishes the approach here from the revelation principle. A standard direct mechanism could be interpreted exactly this way except that any history of messages will generate a probability distribution over profiles of actions.

The set of paths H can be thought of as representing all the histories of the game. At each node along this path some player has an opportunity to send a message. The set of terminal nodes M, represent the set of all possible sequences of messages sent by the players. The mapping λ which maps terminal nodes into outcomes determines the meaning of each of the messages. The only thing unusual about this is the interpretation of the actions in the game since each action is described as a message. The reason for this approach is that some messages are interpreted as commitments. So if a player announces that he is going to have an auction with a certain reserve price, then that is interpreted as a single message or edge in the graph.

Let $P(i)$ be a function that gives the set of nodes where player i sends a message and let $\mathcal{I}_i$ be a partition of $P(i)$ representing what i knows when he sends a message. Each element of $\mathcal{I}_i$ is an information set for i. As competing mechanism games may well have imperfect information as well as incomplete information, player j may not know what commitment message i chose earlier in the game. So this conceptualization captures any model of competing mechanisms in which players have perfect recall.[10]

A behavioral strategy for a player specifies a mixture over his available messages at each information set for each of his types. The player's beliefs at the information set along with his knowledge of others behavioral strategies will allow him to determine the probability distribution over terminal nodes m associated with each of his available messages.

At this point we need some information about the extensive game the players are engaged in to provide a characterization. The difficulty that arises is that the extensive game might impose restrictions on some of the players but not on others. If so, it might not be possible to characterize outcomes in a useful way. To see this, suppose that the default game (the part of the game that is understood) is matching pennies. In the contracting game that lies on top of matching pennies, player 2 is allowed only two different messages. Each of these messages is a commitment message, one of them commits him to heads, the other to tails. Player 1 sends his commitment message at the same time as player 2 - but he has four possible commitment messages instead of just two. One of them commits him to match whatever 2 does, one commits him to do the opposite of player 2, while the other two message commit him to actions that are independent of what player 2 does. If player 1 wants to match player 2, then all equilibrium outcomes will have 1 using a contract that responds to what 2 does. The equilibrium payoffs will be 1 for player 1 and -1 for player 2.

Notice that in this game, player 1 gets strictly more than his min max payoff in the default game. The min-max payoff in the default game tells us nothing about the payoffs that are supportable in the contracting game.

This isn't such a big deal for this special game because the competing mech-

[10]It is somewhat reasonable to imagine that in a complex competing mechanism game players might forget messages they received from other players, or even forget messages that they previously sent. This isn't considered here.

anism game built on the default game only has a single equilibrium outcome. However, we cannot find this equilibrium by knowing only the default game - we need to know which player has access to the bigger set of contracts. Perhaps more important, notice that the revelation principle applied to the default game doesn't characterize the supportable outcomes. The revelation principle has a coordinator who proposes an agreement. If the agreement is refused by one of the players, the coordinator can instruct the non-deviation player to punish the deviator. However, this punishment can't depend on what the deviator actually does. So the revelation principle would say that the coordinator could set up a game to support any outcome in which each player gets at least his min max payoff. However, as we just explained, in the competing mechanism game, the only equilibria have player 2 getting less than his min max.

Of course, we could again use the revelation principle if we could change the default game to one in which player 2 moves first - but that is different from the default game that is observable. The point of this example is simply to illustrate that we can't give any kind of accurate characterization of supportable outcomes from information about the default game alone. We need to know something about the competing mechanism game that is built on top of it.[11]

The way this is handled in [3] and [1] is to assume that the class of games that can be built on top of the default game is very small and consists only of games in which each player simultaneously sends a single message, and in which each profile of messages determines the ultimate probability distribution over actions through a central coordinator. The class of competing mechanism games we consider as possible here is much different. The set of possible messages players can send is much larger. As players may send messages sequentially, we have to worry about sequential rationality. For this reason, we use a different approach based on a concept called *regular games.*

Let $\{\sigma_i, b_i\}_{i=1,...n}$ be behavioral strategies and beliefs for the players specifying mixtures over messages available to players in each of their information sets, and beliefs about the history of play prior to the information set. Let ι be an information set for player i. The continuation game associated with ι is the extensive form game of incomplete information in which each player's type is his payoff type from the original game along with his information about the history of play prior to ι. Beliefs for player i in this continuation game are given by $b_i(\iota)$. For every other player j, the player's type t_j in the continuation game describes (among other things) the most recent information set ι_j in which he

[11] One immediate conjecture that arises is that the problem in this example arises from the fact that player 1 and 2 are asymmetric - they have very different contracting possibilities. One might guess that simply assuming some kind of symmetry in the competing mechanism game would overcome this problem. If both players have the ability to write contracts that depend on the action of the other player, as player 1 does above, then the game doesn't even have an equilibrium, in the sense that if player 2 uses a version of player 1's contract that tries to mismatch player 1's action, then there is no outcome in the game that is compatible with the two contracts. This issue was discussed in [16] who showed that if players are allowed to use contracts that are definable functions of the other player's contract, then not only does this issue does not arise, but a complete characterization of equilibrium outcomes is possible.

sent a message. So player j's belief in the continuation game coincide with his beliefs in the information set ι_j. Using i's beliefs in the information set ι, we write $\rho(t_i, t_{-i}|\sigma_i', \sigma_{-i}, \iota)$ as the outcome function conditional on attaining information set ι when the other players are using the continuation strategies σ_{-i} while player i deviates to continuation strategy σ_i'.

Given a profile of behavioral strategies $\{\sigma_i, \sigma_{-i}\}$, a collection of information sets $\mathcal{I}$ is *attainable with probability* π by player i in the continuation game associated with information set ι if there is a continuation strategy σ_i' for i at ι such that an information set in $\mathcal{I}$ is reached with probability at least π given i's beliefs $b_i(\iota)$ and the continuation strategies σ_{-i}.

An information set ι for player i has the *no-commitment* property if (i) the outcome function $\rho_{A_{-i}}(t_i, t_{-i}|\sigma_i', \sigma_{-i}, \iota)$ is independent of i's continuation strategy σ_i', and (ii) for each $a_i \in A_i$, there is a continuation strategy σ_i' such that $\rho_{A_i}(t_i, t_{-i}|\sigma_i', \sigma_{-i})$ assigns probability 1 to the action a_i. In words, a no-commitment information set is one in which i can carry out any action he likes without changing the behavior of the other players.

Definition 1. *A competing mechanism game is said to be* regular *if for every profile σ of strategies feasible in the game, each player i has for each of his types t_i a strategy σ_i' that attains some no-commitment information set with probability 1.*

This is the extensive form version of the invariant punishment assumption in [16]. The matching pennies example described above obviously isn't a regular competing mechanism game. If player 1 uses the contract that matches player 2's action, then there is no strategy available to 2 in the game in which player 1's action doesn't depend on player 2's action.[12] The intent of the definition should be apparent. We have described the default game as a normal form game and we want to provide a characterization based on the properties of this normal form game since these are the properties that can be verified by what is known to a modeler. This isn't going to be possible if the competing mechanism game completely alters relationship between players that exists in the default game. The regularity assumption is a condition that does this without altering the fundamental richness of the class of competing mechanism games.

As mentioned above, in the models described in [16], and [21] each player offers a contract that specifies a subset (which could consist of a single action) of all his available actions that depends on the Godel encoded value of the other players contracts. One such contract simply specifies that a player reserves the right to select any of his actions in the second stage of the process. Then no matter what contracts the others are using, in the information set in the second stage where the player chooses one of his actions he will be able to select any of his pure actions without changing the behavior of the others. So

[12] In the literature, the recommendation games discussed in [14] and [15] are not regular games. The game between Turing machines discussed in [20] is not a regular game. We discuss these papers below.

every such information set that is attained after offering such a contract has the no-commitment property and the player can attain one of these information sets for sure just by offering this degenerate contract. This effectively provides a Myerson like 'opt out' of the contracting process that supports the usual Individual rationality condition. Definition 1 does the same thing for extensive games.

Digital markets may involve very complicated interactions between players. However, in the end, a buyer's choice consists of the set of prices he or she is willing to accept. Though a buyer could voluntarily make commitments in any information set, we normally imagine these commitment to be voluntary. By simply refusing all commitments until a final purchase choice is made, the buyer can prevent sellers from reacting to her decision rule. Similar remarks apply to sellers. So digital markets plausibly involve regular extensive form games.

Recall the previous example of a non-regular game built on top of the game of matching pennies. In the contracting game, player 1 was able to write a contract to condition his action on the action of player 2, but not conversely. In that game, player 1 starts in an information set that has the no-commitment property. Player 2 only has a single information set. If player 1 chooses the contract that matches actions, then 1's action must depend on the continuation strategy of player 2. So the game does not fit definition 1.

Notice that if we simply viewed the default game as a *sequential* game of matching pennies, this problem goes away since player 1's action choice in the sequential version of the matching pennies game is a rule that specifies a reaction. This reaction function doesn't change as player 2 changes his action. We could then use the sequential default game to characterize the set of supportable outcomes.

What this means is that regularity is an assumption about the relationship between the unobservable contracting game, and the observable default game. It is relatively straightforward to check whether a given contracting game is regular. As the matching pennies example illustrates, it is also easy to see what to do if a particular contracting game isn't regular. The restrictiveness of the assumption lies in its presumption that regularity holds for all the contracting games that might extend the default game.

The first part of our theorem is then given in the following:

Theorem 1. *Suppose the outcome function ω can be supported as a weak Perfect Bayesian equilibrium in some regular competing mechanism game. Then there is a collection of punishments $\{\rho_i\}_{i=1,\ldots,n}$ such that (1), and (2) hold.*

The formal proof is in the appendix, but the logic is straightforward. If the competing mechanism game is regular, then any player can deviate to the strategy that leads to a no-commitment information set. This deviation induces some continuation play that determines the punishment ρ_i. Since this is a deviation from an equilibrium in the game, it must be unprofitable for the deviator no matter what he does in the continuation. This guarantees that (2) holds.

Observe that if the competing mechanism game isn't regular, this is the part of the theorem that will fail, as it does in the matching pennies example discussed above. The reason that the characterization theorem above is useful is because the existential part of the theorem simply makes a statement about a probability distribution over actions that must satisfy a set of inequalities. Since the outcome function is supported as a perfect Bayesian equilibrium, a subset of the set of possible deviations must be unprofitable. If a competing mechanism game isn't regular, then a couple of things go wrong. First, for each deviation we need to worry about what response the other players will make to the deviation. If it were just a matter of picking a different response to each deviation, this wouldn't be such a big deal. In games of complete information, for example, it is a textbook exercise to show that min max and max min payoffs are the same provided non-deviators correlate their actions. What is more troublesome is that each response must itself be part of a continuation equilibrium. So the existential statements are about equilibrium, not about profiles of actions. In a purely computational sense, finding solutions to inequalities is an easy problem, while finding the set of all continuation equilibrium outcomes is a hard problem. Similar problems arise when trying to characterize sequential equilibria of repeated games of incomplete information, for which results exist only for special cases (see for example [22]).

2.2. Reciprocal Contracting

We turn now to the reverse problem. If there is a collection of punishments ρ_i that make an outcome function ω individually rational, we want to show that there is an equilibrium of some regular competing mechanism game that supports ω as a perfect Bayesian equilibrium. It is tempting to conclude that since Theorem 1 already shows that outcome functions can only be supported if they satisfy (1) and (2), then we could just use the generalized revelation principle as stated by [2] to conclude that the outcomes can be implemented with direct mechanisms. With this interpretation, there seems nothing more to prove.

This is unsatisfactory for two reasons. First, it requires an outside mechanism designer to process type reports then instruct each player individually about which action he or she should take. This doesn't seem to fit large digital markets in which mechanism designers only make commitments about their own actions, and only communicate with their own 'agents' - i.e. players who choose to communicate with them. The whole question here is whether these outcomes can be implemented with indirect mechanisms that satisfy these two requirements.

The second difficulty is that the revelation principle is essentially normative. Recall that the problem in a digital market is that the game that participants engage in is only partly understood by the modeler. The fact that whatever this game is can be expressed as a direct mechanism (or a commitment game) doesn't provide any information that can be used to identify this game. The whole point of the revelation principle is to abstract away from models in which detailed messages are sent back and forth between players by putting them into a black box described by a centralized coordinator.

In practice, some of the actual messages that players exchange may be (or may become) observable to a modeler. The reciprocal contracting game that we develop below, uses abstract messages as well. However, since its whole point is to consider mechanisms in which players exchange messages with each other instead of a coordinator, it reveals the nature of the communication that needs to occur to implement specific outcomes. In principle, this should make it possible to determine whether the information transfer that is observable in digital markets could possibly support reciprocal contracting logic.

Finally, the reciprocal contracting game uses the approach in [16] in which contracts can condition directly on other contracts. However, what [16] assume is that players declare publicly a verbatim description of their mechanism - for example, they might publicly post the computer code they use to determine their price. This code is no different in principle from any other message they send, so other players should be able to condition on this message. However, in order to simplify the logic slightly, the reciprocal contracting game assumes that players can describe their commitments 'parametrically'. What this means is that there is a fixed set of messages that players can send at the beginning of the game. Each message in this set describes a commitment. The rules of the default game (which may have been set by some centralized mechanism designer) force the player to obey the commitment associated with any message that he sends.

An analogy is simply the way we typically describe a fixed price mechanism. Literally, the fixed price mechanism names a price then asks the 'agent' to name a quantity. When sent together, the two messages commit the seller to deliver the quantity named by the buyer, and commit the buyer to pay the seller the product of the two messages. Instead of writing all that down, we simply allow the seller to describe his commitment parametrically by naming a price.

The advantage of using this obvious analogy here is that we can write down contracts that are self-referential (in the terminology of [16]) without having to worry about any infinite regress. As will be apparent below, the parametric approach makes it possible to think through the reciprocal contracting game using logic that is almost as simple as the well know Nash Demand Game.

The *reciprocal contracting game* takes place over three stages. In the first stage, each player sends a public message, part of which describes his mechanism, part of which describes an *encrypted type report*; and part of which describes an encrypted number in the interval $[0, 1]$ called a *correlating message*. At the same time each player sends private messages to each of the other players describing encryption keys which can be used to decode the encrypted part of their public message. In the second stage players tell each other what private messages they received at the end of the first stage. All messages from the first two stages are ultimately verifiable and can be used to make commitments. The third stage of the game is a trivial one in which any player who is uncommitted chooses his action in the default game.

Mechanisms are commitments that are based on the various messages players receive during the game. The mechanisms are described by messages in a set $\triangle$ to be defined formally momentarily. As in [16], the commitments associated with the various mechanisms can depend on the public messages that other players send to describe their mechanisms as well as on private messages.

The public messages involve encrypted variables. The treatment here is taken from [15]. Encryption begins with a publicly known function $\kappa : T_i \times [0,1] \times K_1 \times K_2$. The sets K_1 and K_2 contain *encryption keys.* The function κ must have two properties. First for any k_1 and k_2, the function $\kappa(\cdot, k_1, k_2)$ must be one to one and onto some set $\mathcal{E}$. Secondly, for any $k_i \in K_i$ and $e \in \mathcal{E}$, the set $\{(t_i, x) \in (T_i \times [0,1]) : \exists k_j \in K_j; \kappa((t_i, x), k_i, k_j) = e\} = T_i \times [0,1]$. In words, knowing the encrypted value of some variable and one of the keys conveys no information at all about the unencrypted value of the variable. To make the notation a bit simpler, it is assumed that T_i is the same for each player. Then define $\mathcal{E}$ to be the image of the function κ over the domain $T_i \times [0,1] \times K_1 \times K_2$.

The private messages are assumed to come from $K_1 \cup K_2 \equiv K$. Then in the first round, each player sends a public message in $\triangle \times \mathcal{E}$, and a private message in K to each of the other players. At the beginning of the second round, each player will have seen each of the public messages describing commitments, and will have received $n-1$ messages from the other players giving some information about their encryption keys. Each player reports the $(n-1)$ keys he received at the end of the first round to each of the other players. At the end of the second round, each of the other players will have sent a given player n messages. This means each player has $n(n-1)$ private messages from the other players that he must use in some way to figure out their encryption keys.

At the risk of adding notation, just assume that the set of private messages that a player receives is represented by the conventional notation M, where, of course, $M = K^{n(n-1)}$. In our approach, the message space which is processed by a mechanism is $(\triangle^n \times \mathcal{E}^n \times M)$ and a mechanism for player i will be a mapping from each of the points in this message space into his action space A_i. For reasons that will become apparent later, there is no need to consider stochastic mechanisms.

To describe the mechanisms that players are going to use in equilibrium, it is necessary to describe how the private messages are used. Each player is going to commit to a method of translating the private messages in M into a pair of encryption keys for each of the players. The encryption keys will be used to decrypt the public messages in $\mathcal{E}^n$. The decrypted versions of these messages will be treated as type reports (and correlating messages) submitted to a direct mechanism.

The exact method of interpreting the private messages will be borrowed from [15]. Discussion of the details of this conversion can be deferred to an appendix, so for the moment define the mapping $\tau_i : M \to (K_1 \times K_2)^n$ which simply takes a collection of these private messages and uses them to assign a pair of encryption keys to each player. This mapping will represent a part of the commitment that player i makes at the start of the game.

To begin, let $\hat{T} \equiv \prod_j [T_j \times [0,1]]$ be the set containing type declarations and correlating messages of all the players. Let D_i be the set of measurable mappings

$d_i : \hat{T} \to A_i$. We refer to each of these mappings as a *direct mechanism* even though the message space is bigger than the type space. This should create no confusion below. Players are not allowed to offer direct mechanisms (though they can emulate them) in the reciprocal contracting game. Instead, they specify their commitments by using messages in a special set.

The parameter that defines a mechanism is a list $\delta^i = \left\{\delta^i_1, \ldots, \delta^i_n\right\}$ where $\delta^i_j = \left\{d_j, \{p^k_j\}_{k \neq j}\right\} \in D^n_j$. We describe the commitments associated with such a message below. To think about the message that describes a mechanism it might help to think of the analogy with a *proposal.* For example, δ^i_j is a part of the message describing the mechanism used by player i. It could be interpreted to be the direct mechanism that i proposes that player j use, along with the punishments that player i thinks player j should impose when any player k 'deviates'. In this sense, each of the messages describing mechanisms could be thought of as a proposal about how the game should be played by each of the players.

The set of messages $\triangle$ that each player can use to commit himself is then the set of all messages δ^i. Each of these messages is associated with a commitment that translates all the messages the player observes into the action that he takes. The messages he observes are the commitment messages of the other players, referred to here as δ^{-i}; the encrypted public messages $\left(e^i, e^{-i}\right) \in \mathcal{E}^n$; and all the private messages $m^i \in M$ that he receives by the end of the second stage. A mechanism for player i is a mapping from $\triangle^n \times \mathcal{E}^n \times M$ into A_i.

As discussed above, fix a mapping τ from private messages in M to pairs of encryption keys. A precise description of τ will be given below. The commitment associated with the message δ^i is given by

$$\lambda_i\left(\delta^i, \delta^{-i}, \left(e^i, e^{-i}\right), m\right) = \begin{cases} d_i\left(\kappa^{-1}\left(\left(e^i, e^{-i}\right), \tau\left(m^i\right)\right)\right) & \exists \delta^* : \delta^j = \delta^* \forall j \\ p^j_i\left(\kappa^{-1}\left(\left(e^i, e^{-i}\right), \tau\left(m^i\right)\right)\right) & \exists \delta^* : \exists! j : \delta^j \neq \delta^* \wedge j \neq i \\ a_i \in A_i & \text{otherwise.} \end{cases} \tag{3}$$

In this notation, p^i_j and d_i represent the corresponding elements of δ^i and the notation $\exists!$ means "there exists a unique". The function $\kappa^{-1}\left(\left(e^i, e^{-i}\right), \tau\left(m^i\right)\right)$ means that the process τ is used by each player to compute a collection of encryption keys from the private messages. The process that each player uses to create the encryption keys is the same, but whether or not players use the same keys depends on what messages are sent.

One aspect of this that may deserve comment is the fact that when some player j is a unilateral dissenter, the others punish him by using mechanisms that make use of his type and correlating message. Of course, the others can ignore his type and correlating message if they want. If they choose to make use of that information, then the way they use it should be designed so that the dissenter is happy to report his information truthfully.

This formalism now makes it possible to state the main theorem.

Theorem 2. *If there are seven or more players, any outcome function ω satisfying (1), and (2), can be supported as a perfect Bayesian equilibrium outcome in the reciprocal contracting game. Along the equilibrium path of this game, all players make a common proposal δ^*, post encrypted values of their true type along with a correlating message chosen using a uniform distribution on $[0,1]$, 'report' their encryption keys truthfully to the other players, and truthfully report the messages they receive to other players.*

The full proof is contained in Section 4.3. This proof collects results from a number of other papers, so the details aren't new. A sketch of the problems that need to be resolved in the proof, and the method used to resolve them might be helpful.

The first complication stems from the fact that the outcome function being implemented, ω is a joint randomization that may include a lot of correlation between the actions of the players. This must be implemented by independent mechanisms in which players commit only to their own actions in response to messages from other players.[13]

The device that ties the actions of the players together is the correlating messages that players encrypt along with their type reports. The proof adopts a method from [3] (and the generalization provided in [15]) which converts the private correlating messages into something that works like a public randomizing device that the players cannot manipulate. The logic in the proof explains precisely how the mechanisms d_i that players implement along the equilibrium path can be constructed in order to implement the correct joint randomization.

The second complication emerges because players have to have a way to commit to a punishment before they know that some player has deviated. The reason this is important is because once players see that a deviation has occurred, carrying out a punishment may not be sequentially rational. The encrypted types they publish in the first period are there precisely to act as a commitment device. Furthermore, since the players know what deviation has occurred they could provide the others with information that would allow them to carry out a punishment that depends on the deviation. This is what happens in the recommendation game described in [14]. If they can do that, then the competing mechanism game is no longer regular. The set of supportable outcomes should then be larger than the set described here, though there are no formal results about this.

The fine line that the proof has to tread comes from the fact that the first period messages have to indicate that the players have already committed to a punishment without actually sending any information about the types of the players. This is why the messages are encrypted. At the same time, the players have to be able to decrypt these messages in an appropriate way once they learn what the others have offered in the first period. This is where the second round of communication plays a role. The method followed is borrowed from [15].

[13] In other words, they cannot rely on a coordinator whose instructions would correlate their actions.

This is where the requirement that there be 7 players emerges. The reason for the restriction is that the incentive for accurate communication in the second round is created by the belief that the others are communicating accurate information and that a message that disagrees with the messages being sent by others will be ignored. Each player needs to figure out two encryption keys. A player that learns both keys in the first round of communication can determine another player's type, and this will ruin incentives. For this reason each player has to convey each of his encryption keys to two different groups of players in the first round.

In the second round, when a player decides whether or not he should accurately convey the key that he received in the first round, he has to believe that there are at least two other players who have received the same key and that they will report this key truthfully to the other players. This means that in the first round each player must send each of his encryption keys to a distinct group of players, each of which contains at least 3 of the other players. The sender, plus the two distinct groups of 3 players each makes a total of 7 players.

The communication strategies in the second stage depend in a critical way on the outcome of the first stage. If there are more than two distinct proposals made in the first stage, then all players are expected to convey random noise in the second round. Since the message space is a continuum, the probability that two messages agree is zero. The interpretation process τ ignores the second round messages in this case, so it is sequentially rational for all players to send random noise.

The other special case occurs when there is a unilateral deviation in the first stage. In that event all players are expected to send random noise to the deviator in the second round, but accurate messages to the non-deviators. Once again, this behavior is self enforcing when all the other players are expected to follow it.

When there are six or fewer players, the construction in the reciprocal contracting game breaks down. Players can no longer guarantee that their type will be accurately revealed to non-deviating players. Type contingent punishments are then more difficult to enforce. At least in the context of the reciprocal contracting game, the set of outcomes that can be supported as equilibrium will shrink. A counterexample where this matters would evidently require that the preferences of the deviating principle depend on the actions of the non-deviators in a way that depends on their types, so preferences would be interdependent. A counterexample would likely be complex and abstract. Ultimately the counterexample wouldn't address the real issue anyway, since there may be games other than the reciprocal contracting game that support the problematic allocation. For this reason we leave the study of this problem in games with fewer than six players to another paper.

3. The Equivalence of Competing Mechanisms and Reciprocal Contracting

Combining Theorem 1 with Theorem 2 gives the following corollary:

Theorem 3. *An outcome function ω is supportable as an equilibrium in a regular competing mechanism game with seven or more players, if and only if it is supportable as an equilibrium in the reciprocal contracting game.*

Recall that the reciprocal contracting game is intended as a default game. Players could try to devise ways around outcomes that they didn't like by offering their own contracts outside the game. What Theorem 3 shows is that even if they did this, no new equilibrium outcomes could be supported. In this sense, the reciprocal contracting game is robust to the invention of new trading methods.

One implication of Theorems 1 and 2 is the not wholly surprising result that if one doesn't know exactly what contracting game players are engaged in, then it will be hard to predict their behavior. What is surprising here is that despite this, the set of outcome function that can be supported as equilibrium can be characterized provided one does know that the game the players are playing is regular.

This still leaves a large set of equilibrium outcomes. One natural way to refine them would be to assume that certain types of communication aren't possible, or that certain types of commitment aren't feasible. An example of this approach is provided by [21]. He splits the players into informed agents and uninformed principals and assumes that agents can only commit themselves to an interaction with a single principal (exclusive agency). This is similar to the approach in the early literature on competing auctions [23] or [24] where uninformed sellers acted as principals making commitments to informed buyers who would commit only to the auction they bid in.

A very different approach is taken in the paper by [25], who uses the concept of strong robustness from [4] to refine the set of reciprocal contracting equilibria. Strong robustness requires that when a mechanism designer changes his mechanism in a competing mechanism game, the deviation remains unprofitable even if the deviating mechanism designer could pick the continuation equilibrium that follows the deviation. [25] shows that in a large auction market, the competitive allocation can be supported by equilibria that are strongly robust, and conjectures that this is the only allocation that has this property.

Finally, a lifetime of doing mechanism design makes the Myerson coordinator seem natural. The coordinator just asks all the players to report their types then tells them what to do. In an auction environment, there is a coordinator - the seller. In a market environment it is less clear that such coordination exists. On one hand there has to be coordination to set the legal rules associated with contracts and to provide the conventions that decide the rules of the contracting game that players actually play. This paper imposes restrictions of what the coordinator can actually accomplish by requiring that all enforceable contracts commit only the actions of the player who offers the contract. One way of interpreting Theorems 1 and 1 is simply that there is no conflict between these two approaches provided one is confident that the mostly unobservable market game is regular. The theorem illustrates how the market can implement the correlation and communication that the centralized coordinator is supposed to provide. In many ways, this summarizes the message of the paper - one can use Myerson to understand digital markets provided one is confident the market game is regular.

Conclusions

Theorem 2 characterizes the set of outcome functions that can be supported as equilibrium outcomes in regular competing mechanism games. Theorem 3 shows that all these outcome functions can be understood using a single reciprocal contracting game. In this sense reciprocal contracts play the role of direct mechanisms for competing mechanism games, but in a specially simple way since only a single contracting game is required to understand all the supportable outcomes.

To summarize the results in as short a way as possible, the set of outcome functions supportable as equilibrium in competing mechanism games in general and digital markets in particular is the set described in Myerson's textbook. Does the fact that competing mechanism games are extensive form games while Myerson is about normal form games make any difference? No. Does the fact that competing mechanism games use a different solution concept than Bayesian equilibrium make a difference? No. Does the fact that competing mechanism games involve players making their own commitments rather than submitting to a coordinator make a difference? No. Do we need a collection of games to understand these outcome the way we need a collection of direct mechanisms to understand mechanism design? No, one game will work for all outcomes.

What does matter? The paper suggests two things. First, it probably isn't particularly useful to propose specific extensive form games that seem descriptively reasonable, at least for digital markets. The reason is that commitments in actual markets have become (or are becoming) too complex and hard to observe to model descriptively. Secondly, if the set of outcome functions described here seems too large, then some reason has to be provided to explain why firms can't use reciprocal contracts.

It isn't particularly hard to see why firms might not be able to use reciprocal contracts - they involve complex commitments and communications that seem implausible. Saying that doesn't particularly help in providing broad stroke restrictions on contracts that are convincingly immune to the kinds of innovation that goes on in digital markets so that they might help in improving the predictive power of competing mechanisms. This paper doesn't have much to add on what these broad stroke restrictions might be. Hopefully the theorems here at least narrow the search by ruling out some of the obvious possibilities.

4. Appendix: Proofs

4.1. Proof of Theorem 1

Proof. Let $\omega(t)$ be the outcome function supported by some equilibrium of a regular competing mechanism game in which strategies are σ^*. It satisfies (1) by the usual revelation principle.

The approach in the proof is to pick a out of equilibrium behavioral strategy. Deviating to this strategy will be unprofitable. The punishment is going to be the response the other players make to this deviation. The problem is to show that this strategy can be chosen to have the properties described in Theorem 1.

The game is regular, so i has a behavioral strategy σ_i' that attains a no-commitment information set with probability 1. Since the set of strategies available to players doesn't depend on their type, we can choose this strategy so that the collection of no-commitment information sets it induces is the same for every one of the player's types. Write $\mathcal{I}\left(\sigma', \sigma_{-i}^*\right)$ be the collection of no-commitment information sets in the support of the strategy pair $\left(\sigma_i', \sigma_{-i}^*\right)$. By definition, each no-commitment information set has the property that there is a collection of strategies that ensure that no matter what pure action the player wants to play in the continuation, he can play that action and induce the same reaction from his opponents. Since this reaction can't depend on the deviating player's type, write it as $\rho_{A_{-i}}\left(t_{-i}|\iota\right)$.

The deviation we want to focus on is the one in which the player adopts some behavioral strategy that attains a no-commitment information set with probability 1. Then whenever the player attains such an information set, he chooses his favorite action conditional on updated beliefs, then implements it using the strategy that induces the common response $\rho_{A_{-i}}\left(t_{-i}|\iota\right)$. With a slight abuse of notation, call this strategy σ_i'.

Of course, deviating to this strategy is unprofitable because σ_i^* is part of a perfect Bayesian equilibrium. The payoff function that prevails when player i uses this behavioral strategy σ_i' can be written

$$\mathbb{E}\left\{u_i\left(\left(\sigma_i', \sigma_{-i}^*\right), t\right)|t_i\right\} =$$

$$\mathbb{E}_{\iota\in\mathcal{I}\left(\sigma_i', \sigma_i^*\right)}\left\{\max_{-a_i\in A_i}\mathbb{E}\left\{u_i\left(a_i, \rho_{A_{-i}}\left(t_{-i}|\iota\right), t\right)|t_i, \iota\right\}\right\}$$

by the law of iterated expectations.

By Blackwell's Theorem, this expression is no smaller than

$$\max_{-a_i\in A_i}\mathbb{E}_{\iota\in\mathcal{I}\left(\sigma_i', \sigma_i^*\right)}\left\{\mathbb{E}\left\{u_i\left(a_i, \rho_{A_{-i}}\left(t_{-i}|\iota\right), t\right)|t_i, \iota\right\}\right\} =$$

$$\mathbb{E}\left\{u_i\left(a_i, \mathbb{E}_{\iota\in\mathcal{I}\left(\sigma_i', \sigma_i^*\right)}\rho_{A_{-i}}\left(t_{-i}|\iota\right), t\right)|t_i\right\}$$

by the linearity of preferences. This verifies that the collection of punishments $\left\{\mathbb{E}_{\iota\in\mathcal{I}\left(\sigma_i', \sigma_i^*\right)}\rho_{A_{-i}}\left(t_{-i}|\iota\right)\right\}_{i=1,\ldots,n}$ satisfies (2). □

4.2. The transformation τ

The most complex part of the reciprocal contracting game is the transformation τ that is used to generate encryption keys from the collection of private messages m that a player receives. The details of this transformation are described formally in [15], so we provide a heuristic description here.

Focus on some player i, and imagine that player i chooses two distinct groups of friends F_i^+ and F_i^- from among the other players. The groups F_i^+ and F_i^- are mutually disjoint, and each group consists of at least three of the others. Player i and his two groups of three friends each explains the restriction in the theorem to 7 or more players.

In the first round, player i is going to report his first encryption key to his friends in F_i^+ and his second encryption key to his friends in F_i^-. The single encryption keys don't convey any useful information to his friends about his type - they need both keys to decode his public messages. Player i's friends are commonly known by all players. Every other player does the same thing - chooses two distinct groups of friends and reveals one of his encryption keys to each group of friends privately in the first round.

In the second round, every player can try to learn i's encryption keys by asking his friends about them privately. There is nothing that compels his friends to reveal the keys accurately, so the other players have to interpret the friends messages in a special way in order for the information to be revealed.

The way this is done is to imagine that when player i receives reports from i's friends, he ignores the reports and uses some arbitrary encryption key, unless at least two of these reports agree. If two or more of the reports agree, then player j accepts this majority report as i's actual encryption key.

The reason this works is that if each of i's friends believes that i reported his encryption key accurately to each of them, and believes that each of i's other friends are going to convey those reports accurately to other players, then there is nothing they can do to manipulate the interpretation of the reports since their contrarian messages will be ignored. It is therefore always sequentially rational for them to report the encryption key they heard accurately, no matter how the information is ultimately used.

Player i could try to manipulate this by sending different encryption keys to his friends, but if the outcome function being implemented is incentive compatible, he cannot profit by doing so.

A formal description of this process is complicated by two small details. First, the transformation τ is supposed to be the same for every player. To accomplish this, the reciprocal contracting game assumes that i commits himself to use the encryption keys his friends report to him to decipher his own public message from the first round, forcing him to interpret the public message the same way that other players do.

The second complication occurs when player i is a friend of some player j whose encryption key he needs to decrypt. In that case, i commits to use the reports he hears from j's other two friends along with the private message he received from j in the first period.

To provide a formal description of τ, let m be the vector of $n\,(n-1)$ messages that i receives from the other players, as described above. Let $m^j = \left\{m_1^j, m_2^j\right\}$ be the report that i receives from player j. The first component of m^j is the private message that i receives from j in the first round. The message m_2^i is a vector of $(n-1)$ messages that player j heard at the end of the first round.

The notation m_{2k}^{j} is the encryption key that was reported to player j by player k at the end of the first stage of the reciprocal contracting game. Define

$$\mathcal{M}_j^+ (m) = \begin{cases} m' & m_{2j}^k = m' \forall k \in F_j^+ \vee \exists ! m_{2j}^k \neq m'; k \in F_j^+ \\ \emptyset & 0\text{therwise.} \end{cases} \tag{4}$$

The notation $\mathcal{M}_j^+ (m)$ means the majority report of j's friends in F_j^+. The idea is that if j's friends in F_j^+ (to whom j is supposed to report his first encryption key) all agree about the key they were given, or if only one of them disagrees, then i should use the key they reported to decrypt j's message. To compute the first of j's encryption keys, i would then use the formula

$$\tau_j^1 (m) = \begin{cases} \overline{\kappa} & \mathcal{M}_J^+ (m) = \emptyset \\ \mathcal{M}_j^+ (m) & \text{otherwise.} \end{cases}$$

A similar formula describes the encryption keys that i commits to use for each of the other players.

Notice that the transformationτ is part of player i's mechanism. He commits to use it to decrypt the public messages. As such, the convention below is to adopt this transformation as part of the reciprocal contracting game

With this last construction, it is possible to state and prove the main theorem of the paper.

4.3. Proof of Theorem 2

Proof. Begin with the equilibrium path

Index the action profiles in A in some arbitrary way. Let $\omega^k (t)$ be the probability assigned to action profile a^k by the outcome function ω when player types are given by the vector t. The notation a_i^k means the action taken by player i in action profile a^k. For $(t, x) \in \hat{T}_i$, let t_i and x_i be the type and correlating message declared by player i in the first round. The following mapping defines a direct mechanism:

$$d_i^{\omega} (t, x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \omega^{\tau} (t_i, t_{-i}) \geq \lfloor x_i + \sum_{j \neq i} x_j \rfloor \right\} \tag{5}$$

The notation $\lfloor y \rfloor$ means the fractional part of the real number y. This function aggregates the correlating messages into a number between 0 and 1, then uses this to choose an action profile in A. The mechanism then commits i to carry out his part a_i^k of the corresponding action profile a^k.

This will implement outcome a^k with probability $\omega^k (t_i, t_{-i})$ provided each of the x_j are uniformly distributed on $[0, 1]$. The property of this construction that will be especially useful below, is the fact that as long as each of the other players is choosing x_j uniformly, the random variable $\lfloor x_i + \sum_{j \neq i} x_j \rfloor$ is uniform on $[0, 1]$

for each value of x_i.[14] What this means is that the probability distribution over i's actions is independent of x_i. As a consequence, it is sequentially rational for i to choose his correlating message uniformly from $[0,1]$ (no matter what his commitment) provided he thinks the other players are doing the same.

By (2), there is a collection of punishments $\{\rho_i\}_{i=1,\ldots,n}$ associated with ω. For every such punishment, define for each of the players other than i the direct mechanism

$$p_j^{\rho_i}(t,x) = \left\{ a_i^k : k = \min_{k'} : \sum_{\tau=1}^{k'} \rho^\tau(t) \geq \lfloor \sum_i x_i \rfloor \right\}. \tag{6}$$

As above, this will implement j's part of the punishment providing reports are truthful and correlating messages are uniform.

We are now ready to give the strategies associated with the Perfect Bayesian equilibrium that supports ω. Each player should propose

$$\delta^* = \left\{ d_i^\omega, \{p_i^{\rho_j}\}_{j \neq i} \right\}_{i=1,\ldots,n}$$

then punish an encrypted version of his true type along with an encrypted value of a number chosen uniformly from the interval $[0,1]$. Each player should then privately give the correct encryption keys to his friends.

In the second stage, if all players offered δ^*, then each player should report the encryption keys he heard at the end of the first round accurately. If a single player j deviated in the first round, then each player should send a completely random set of reports to j, but report the keys they heard in the first period accurately to each of the other players. If there are more than two distinct proposals offered in the first round, each player should send noise to each of the other players in the second stage, then play their part of some Bayesian Nash equilibrium in the final stage.

It remains to check deviations. The 'part of a Bayesian Nash equilibrium' part described above is straightforward, so attention is restricted to the case where all players offer the same mechanism, and to the case where a single player deviates.

If all players offer the same mechanism, and players otherwise follow the equilibrium strategies, then by the argument following (5), each player's payoff is the one associated with the outcome function ω. Since ω is incentive compatible, no player has an incentive misrepresent his type, or send out false encryption keys since his payoff is at least as large when he conveys his true type as it is when this type is replaced by something else when mechanisms are implemented. In the second stage, each player believes the others have conveyed their types truthfully and reported the same encryption keys to each of their friends. So any player who lies about the encryption keys he learned in the first period

[14] This device is from the paper [3] who do this for two players. A proof of this last property when there are more than two players is given in [15].

believes that his report will be ignored (see (4)). In this history, no player has a move to make in the final stage, so the continuation strategies are sequentially rational in any history in which all players offer the same mechanism.

In a history where there is a single deviator, the argument is similar for the non-deviators. They have no incentive to lie to the other non-deviators about the messages they heard in the first stage. The only exception is that each player is supposed to send an uninformative message to the deviator in the second stage. Since the others are believed to be sending uninformative messages, the probability that two messages match is zero. So no player believes his message to the deviator will change any outcome.

As for the deviator, he has no incentive to lie about the encryption keys he learned in the first period since he thinks deviating messages will be ignored. Notice that since the deviator is required to report these messages truthfully to incentivize the others, the deviator participates in his own punishment. The deviator believes the messages he receives about other players encryption keys is uninformative, so the best he can do is to maximize his expected payoff given his interim beliefs against the punishment ρ_j that he expects the others to implement (again following (6)). Since the outcome function ω is individually rational, this deviation is unprofitable. □

References

[1] F. Forges, A folk theorem for bayesian games with commitment, Games and Economic Behavior 78 (2012) 64–71.

[2] R. Myerson, Game Theory The Analysis of Conflict, Harvard University Press, 1997.

[3] A. Kalai, E. Kalai, E. Lehrer, D. Samet, A commitment folk theorem, Games and Economic Behavior 69 (1) (2010) 127–137.

[4] L. Epstein, M. Peters, A revelation principle for competing mechanisms, Journal of Economic Theory 88 (1) (1999) 119–160.

[5] A. Pavan, G. Calzolari, A markovian revelation principle for common agency games, northwestern University Discussion Paper (2001).

[6] A. Pavan, G. Calzolari, Sequential contracting with multiple principals, Journal of Economic Theory 50 (2009) 503–531.

[7] A. Pavan, G. Calzolari, On the use of menus in sequential common agency, Games and Economic Behavior 64 (2008) 329–334.

[8] D. Martimort, L. Stole, The revelation and taxation principles in common agency games, Econometrica 70 (4) (2002) 1659–1673.

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

[10] S. Han, Menu theorems for bilateral contracting, Journal of Economic Theory 131 (1) (2006) 157–178, available at http://ideas.repec.org/a/eee/jetheo/v131y2006i1p157-178.html.

[11] B. D. Bernheim, M. D. Whinston, Common agency, Econometrica 54 (4) (1986) 923–942.

[12] D. Bernheim, M. Whinston, Menu auctions, resource allocations and economics influence, Quarterly Journal of Economics 101 (1986) 1–31.

[13] D. Martimort, Exclusive dealing, common agency and multiprincipals incentive theory, Rand Journal of Economics 27 (1996) 1–31.

[14] T. Yamashita, Mechanism games with multiple principals and three or more agents, Econometrica 78 (2) (2010) 791–801.

[15] M. Peters, C. Troncoso-Valverde, A folk theorem for competing mechanisms, Journal of Economic Theory 148 (3) (2013) 953 – 973. doi:10.1016/j.jet.2013.01.001.
URL http://www.sciencedirect.com/science/article/pii/S0022053113000355

[16] M. Peters, B. Szentes, Definable and contractible contracts, Econometrica 80 (1) (2012) 363–411.
URL http://ideas.repec.org/a/ecm/emetrp/v80y2012i1p363-411.html

[17] I. Segal, M. Whinston, Robust predictions for bilateral contracting with externalities, Econometrica 71 (3) (2003) 757 – 791.

[18] V. Duquietd, D. Martimort, Mechanism design with bilatral contracting, Tech. rep., Paris School of Economics Discussion Paper (2012).

[19] M. J. Osborne, A. Rubinstein, A Course in Game Theory, MIT PRess, Cambridge Massachusetts, 1994.

[20] L. Fortnow, Program equilibria and discounted computation time, Working paper, University of Chicago (january 2009).
URL http://www.cs.uchicago,edu/fortnow/papers/discount.pdf

[21] B. Szentes, Contractible contracts in common agency problems, Tech. rep., London School of Economis (2013).

[22] M. Peski, Repeated games with incomplete information and discounting, Tech. rep., University of Toronto, Department of Economics (2013).

[23] P. McAfee, Mechanism design by competing sellers, Econometrica 61 (6) (1993) 1281–1312.

[24] M. Peters, Equilibrium mechanisms in a decentralized market, Journal of Economic Theory 64 (2) (1994) 390–423, manuscript University of Toronto to appear in Journal of Economic Theory.

[25] S. Han, Robust competitive auctions: A theory of stable markets, unpublished, University of McMaster working paper. (July 2014).