COMPETING MECHANISMS

MICHAEL PETERS, UNIVERSITY OF BRITISH COLUMBIA

ABSTRACT. The recent literature on competing mechanisms has devoted a lot of effort at understanding a very complex and abstract issue. In particular, an agent's type in a competitive environment is hard to conceptualize because it depends on information the agent has about what is going on in the rest of the market. This paper explains why this is such an important practical problem and illustrates how the literature has 'solved' it.

There is a very subtle and ingenious technique called *click stream pricing* that many websites use when they make price offers to consumers. You can probably find an example pretty easily for yourself, however the last time I tried I found at quote for a desktop computer on Amazon. The price was \$689 (with free shipping). You can see offer (which probably looks like an offer you have seen many times before). From the point of view of this paper, the interesting thing about this figure is the little check box below the "Add to Cart" button which offers to add a 2 year warranty for my computer to the cart at a price of \$77.99. Checking it, then clicking on the cart button brought up a breakdown of the price, including the 2-year warranty.

Thanks to Li Hao for his comments.

Nothing forces you to check the 2-Year Warranty button. Leaving it unchecked, then clicking on the cart brought up another familiar image from the internet.

Notice the 2-Year Warranty, now at a price of $56.99. Of course, it could be that a "warranty" is different from a "Canopy protection plan". Even if they were the case, there is nothing in the information given to suggest that a Canopy protection plan is worth less (or more) than a warranty.

There are a number of things to observe about this practice. First, people who buy the bundle of a computer with a warranty will pay different prices depending on what they click on to receive the price offer. So much for the idea that commodities like computers, or warranties have well defined prices.

In mechanism design, we would normally think of the clicks that a buyer sends to the company's servers as *messages.* The computer programs that process the messages and sell the goods we would think of as *indirect mechanisms*. For the rest of this paper it might be helpful to think of the programs that drive these websites whenever the word 'mechanism' comes up.

There is another aspect to this process that is probably more important from the point of view of this paper. The reason that I bypassed the initial offer for the warranty to look for the cheaper price had nothing to do with how much utility I would get from a computer or warranty. I clicked through because I had seen the click stream method before. The message I was sending to the seller was about my *market information,* not about my use value.[1]

The reason that this might make a difference to Amazon is that they are not the only supplier of desktop computers and warranties. The mechanism is competing with similar mechanisms at other stores. My message indicated to them that these other stores were offering lower prices than their initial offer of $77.99. This gives Amazon's dumb mechanism an opportunity to react to this 'deviation' and cut their own price. This sounds like a very collusive technique, though it is being used by a seller who is participating in a market that we would normally think of as being very competitive.

This is the concern of this paper: Competition in mechanisms leads to outcomes that behave quite differently from what we might expect even when a 'market' has lots of sellers and there are low barriers to entry.

Of course, most micro economists now understand the basics of mechanism design - buyers report types to a mechanism designer who commits to incentive compatible outcomes. Competition seems to involve little more than having two or more mechanisms that buyers can choose from, then analyzing the buyers' participation decisions.

The click stream pricing example illustrates a sort of basic flaw to this idea. If buyers types contain both payoff information and market information, then buyers' types themselves are endogenous in ways that seem hard to understand. The techniques of mechanism design are elegant, but they don't work at all when buyers' types are endogenous.

The objective of this paper is to illustrate how to understand competition in digital markets using the theory of competing mechanisms.

There are two main conclusions to be drawn from this literature. The first is that there is nothing analogous to the *First Welfare Theorem* in a digital market. Very competitive digital markets might work well, but they can also support undesirable collusive outcomes. Indeed, the main characterization theorem of this literature shows that this is true independent of the underlying allocation of property rights or of the existing institutional setup in these markets.

For this very reason, there is a second implication of this literature. There is little reason to set up careful extensive form models of competition in digital markets in which sellers are restricted to a special set of mechanisms and buyers can only send very specialized messages like

[1]One reason I might click through the initial offer is because I have no use value for the warranty at all. An equally plausible reason for clicking through is that the buyer knows he can buy a warranty for a lower price elsewhere. The point being made is that signals sent by buyers may contain market information.

bids. Since sellers can work around these restrictions using methods that may be hard to understand conceptually, the predictions of these models will be of little use. A more useful way to proceed is simply to give up on the idea that predictive models of competition can be devised, and instead study the set of supportable outcome functions directly in order to get clues about how to make these markets work.

The revelation principle still works, but not as you expect. These examples illustrate a number of things about competition in digital markets. Old fashion textbook models of demand and supply completely miss the point. Goods don't have prices that can clear markets. Instead, transfers that occur when there is a trade are determined by a sometimes complex interchange of messages between buyers and sellers. These transfers are unique to the buyer who makes the trade, not to the good being traded. The rules that websites use to determine price offers more closely resemble *mechanisms* than traditional *markets.* Yet even if you accept that, these mechanisms behave in a fundamentally different way in a competitive environment than they do in standard applications of the revelation principle.

First, the idea that a buyer has a 'type' needs to be re-thought. As in the click stream pricing example given above, a buyer's willingness to pay for something depends not only on his preferences, but also on what he or she knows about prices being offered in the rest of the market. To the extent that these prices reflect some kind of equilibrium, what that means is that buyers' types are actually endogenous.

The idea that buyer types should be redefined to reflect their market information was originally suggested (as far as I know) by Preston McAfee ((McAfee 1993)). He didn't do anything with the idea, beyond acknowledging it to be a problem. The idea was formalized in (Epstein and Peters 1999). What they did was to show that competition in mechanisms would lead to equilibrium outcomes which could be represented in the usual revelation principle fashion with each principal offering a direct mechanism and every agent reporting his or her type to the principal. Agent types reflected all of the agents' market information in exactly the fashion suggested by McAfee. The first major accomplishment in that paper was to show that these types could actually be described mathematically.[2]

[2]What makes this challenging is that there is a sort of infinite regress associated with this. When an agent tries to describe the mechanism being used by another principal, one of the things he has to describe is whether that mechanism chooses an outcome that depends on the mechanism of the principal he is reporting to. If it does, then he must also describe whether the other principal's mechanism chooses

The description of a type was modeled around the idea of willingness to pay that is so widely used in auction theory. When a bidder in an auction describes his type all he has to do is to describe the maximum amount he is willing to pay for a good. When he conveys this information, he describes not only whether or not he will buy at the seller's current price, but also what he might do at other prices. In a competitive environment, this willingness to pay depends on the payoff the buyer can attain by trying to trade with one of the other sellers. This is easy enough to describe when the seller's price is fixed. Yet if the seller changes his price, participation decisions will change and buyers' payoffs with the other sellers will change in complex ways. For the seller to figure out a buyer type, he has to understand how this outside option would change as he varies his price.

(Epstein and Peters 1999) modeled this by imagining that the buyer would answer a series of questions about payoffs under different circumstances. The argument is demanding, however it is possible to understand its basics by imagining a trivial example in which there are two sellers. Seller 1 has two mechanisms a and b, while seller 2 has three mechanisms A, B, and C. Seller 1 wants to ask an agent which of three mechanisms A or B or C that seller 2 is using. Suppose that from the agen't perspective, the mechanisms look as follows:

	A	B	C
a	$\overline{\pi}$	$\pi''' < \overline{\pi}$	$\pi' < \underline{\pi}$
b	$\pi'' < \overline{\pi}$	$\overline{\pi}$	$\underline{\pi} < \overline{\pi}$

Along the top of the table are the mechanisms that seller 2 could offer (seller 1 is trying to figure out which one seller 2 actually used). The rows each correspond to a mechanism that seller 1 could offer. From the inequalities in the diagram, you can see that the agent wants either the combination (a, A) of mechanisms, or the combination (b, B) of mechanisms.

For each of the mechanisms that seller 2 could offer, there is a maximum payoff that the buyer could attain. This payoff is $\overline{\pi}$ if seller 2 is using either mechanism A or B, and $\underline{\pi}$ if seller 2 is using C as is easily read from the table. Call these first order payoffs.

For each of the first order payoffs, there is a maximum payoff the buyer could attain if seller 1 were offering a mechanism that supported the first order payoff. For example, both mechanisms a and b will support the first order payoff $\overline{\pi}$ if seller 2 is offering the right mechanism.

an outcome that depends on whether the first principal's mechanism depends on the other principal's mechanism, etc.

However, the first order payoff $\underline{\pi}$ is only attained when seller 1 uses mechanism b.

This gives a pair of *second order payoffs.* The best payoff the buyer can attain when seller 1 is using a mechanism that supports the first order payoff $\overline{\pi}$ is $\overline{\pi}$ for mechanisms A and B, and $\underline{\pi}$ if seller 2 is using mechanism C. Not much news there. However the best payoff the buyer can attain when seller 1 is using a mechanism that supports first order payoff $\underline{\pi}$ (i.e., mechanism b) is π'' when seller 2 uses A, $\overline{\pi}$ when seller 2 uses mechanism B, and $\underline{\pi}$ when seller 2 uses mechanism C.

We can summarize this in a table:

	First Order Payoff	Second Order Payoffs
Seller 2's mechanism is A	$\overline{\pi}$	$(\overline{\pi}, \pi'')$
Seller 2's mechanism is B	$\overline{\pi}$	$(\overline{\pi}, \overline{\pi})$
Seller 2's mechanism is C	$\underline{\pi}$	$(\underline{\pi}, \underline{\pi})$

Notice now that each of seller 2's mechanisms is uniquely identified by a sequence - for example, mechanism B is associated with the sequence $(\overline{\pi}, (\overline{\pi}, \overline{\pi}))$. So when the buyer is trying to tell seller what seller 2 is doing, he can report this sequence to the seller and it uniquely identifies mechanism B.

These sequences accomplish a couple of things. First, they are expressed in a language that is independent of the messages associated with the mechanisms the sellers offer and the actions that the sellers take. The sequences involve payoffs, which don't change with the message spaces used.

Since the buyer is describing his type to seller 1 by describing payoffs conditional on seller 2's mechanism, it might seem that it would be more straightforward for seller 1 to simply ask the buyer what his payoff is in seller 2's mechanism, rather than dragging out the whole sequence. If Seller 1 were to rely only on the buyer's payoff with seller 2, he faces an intractible complication - if seller 2's mechanism depends on what seller 1 does, then the buyer's type also depends on seller 1's mechanism. If that is the case, the whole apparatus of mechanism design disappears. The second advantage of the approach just described is that by asking the buyer to answer questions about the maximum attainable payoff, the buyer's answers no longer depend on what seller 1 is actually doing.

In (Epstein and Peters 1999) these payoff sequences were referred to as *universal types.* When mechanisms asked agents to reveal their market information by reporting their universal types, they were said to belong to the *universal set of mechanisms.*

Yet the fact that these 'types' are universal is just a part of the contribution. Since the buyer now has a way to report the mechanism seller 2 is using, seller 1 can build in a punishment for seller 2 whenever seller 2 deviates from some putative equilibrium strategy. Provided seller 1 can provide the right incentives for the buyer to report these deviations truthfully, the seller can change a static competing mechanism game into something that works more like a dynamic repeated game.

The messages are way too complicated. In a sense, the message in (Epstein and Peters 1999) is that standard mechanism design and the revelation principle aren't really helpful when thinking about competing mechanism games. The messages that agents need to send to report their types are too complicated to be used. Even if you could use them, neither of the properties that make mechanism design work[3] actually hold in competing mechanism games. It is hard to describe something that is incentive compatible and individually rational in a competing mechanism game because both terms are effectively meaningless when types and outside options are both determined by the outcome of the game. The entire description of equilibrium has become completely circular.

It is important not to confound the theoretical and practical issues in this. Sellers clearly do want information about their competitors actions and clearly do ask buyers to report this information when they use methods like click stream pricing. We would normally call these methods *indirect mechanisms.* An auction is an indirect mechanism. Auctions have the fortunate property they are widely used and easy to recognize in practice. It is also easy to understand the properties of auctions using the revelation principle. The corresponding indirect mechanisms in competitive digital markets are considerably harder to recognize and vary widely across different industries. The actual negotiation of prices is carried out by computer programs that run in some kind of ether that is difficult to observe or understand for anyone but the engineer who programmed the software. That is why it is so critical to be able to rely on something like a revelation principle which can abstract from these coding details in order to understand how the markets are working. The theoretical arguments in (Epstein and Peters 1999) show that the revelation principle as we usually think of it works in an abstract theoretical sense, but not in a practical computational sense.

[3] That is, exogenously distributed types, and exogenous outside options.

What seemed to be needed was some other kind of mechanism that didn't rely on buyers reporting their types. The common agency literature provided the first clue about what this mechanism might be with something called a *menu theorem* ((Peters 2001), (Martimort and Stole 2002), or (Pavan and Calzolari 2009)). This is a generalization of similar result from standard mechanism design that is sometimes referred to as the *taxation principle.* The idea is easy enough in standard mechanism design. When a principal designs a direct mechanism he asks his agent to report his type, then commits himself to an outcome for each possible type that the agent reports. If the mechanism is incentive compatible, the agent will want to report his type truthfully because the outcome associated with reporting truthfully is more desirable than any of the others. So the principal might just as well offer his agent the menu of outcomes associated his direct mechanism and ask the agent which one he wants - the final result will be the same. That is the taxation principle.

Common agency means many principals all compete to influence one single agent. The menu theorems I mentioned above showed that if you model the common agency in whatever way you like, allowing each of the principals to ask the agent to send messages of any kind, and you can find an equilibrium for your model, then you could have supported the same equilibrium outcome by having each of the principals offer the single agent a menu of outcomes that he could choose from. The significance of this is that there are no messages, the common agent simply makes a choice from a set of outcomes involving things like prices that are easily observable. So even if the competition among the principals is carried out through complex computer programs, it should still be possible to understand the competition simply by looking at the outcomes it supports.

Eliminating the complex messages about type is a big advance. I won't go through the details of the menu theorems because they had two troubling flaws. First, though menus are conceptually simpler than direct mechanisms, the literature on common agency didn't explain exactly what kinds of outcomes these menus would support. The best you could do with the menu theorem was to try to find a way to manipulate them in whatever applied problem you happen to be working on. Furthermore, the menu theorem only holds when there is a single agent, which seems to rule out using this theorem to explain digital markets where the are obviously a huge number of agents.

How to use menus with many agents. The way around the multiple agent problem was suggested by Takuro Yamashita (Yamashita

2010) who was a graduate student at Stanford at the time, and is now at Toulouse. His idea was very simple. I'll illustrate it with a complete information example. In this example there are 5 players. The first two of the players are competing principals, the other 3 are agents that the principals interact with. The following table describes the payoff in the game they play:

1\2	Left	Right
Up	$2, 2, 2$	$-1, 3, 0$
Down	$3, -1, 0$	$0, 0, 0$

Principal 1 can play either Up or Down, principal 2 either Left or Right. The payoffs in each cell are listed principal 1 first, principal 2 second, and each of the three agents third. In this example, the agents don't control any actions, they exist simply to send messages to the principals. The agents do care about the outcome, they want the principals to coordinate on the action profile (Up, Left), which is not an equilibrium for the principals left to their own accord.

Here is how Takuro thought this game should be played. First, each principal would design a mechanism and commit himself to follow it. This means perhaps writing a computer program which will carry out some action based on inputs from the agents. The message an agent is allowed to send to principal 1 can be either Up or Down, while an agent can send either message Left or Right to principal 2. Each principal commits himself to implement the action that is chosen by the majority of the agents. In this competing mechanism game, the outcomes (Up,Left) and (Down, Right) can be supported as equilibrium outcomes (note this is all the outcomes in which each principal gets at least his minmax payoff).

To support the outcome (Up, Left) for example, each principal adopts the mechanism that commits to take the action chosen by the majority of the agents. The agents use the following actions in the continuation: if both principals offer this majority mechanism, all three of the agents choose Up from principal 1's mechanism and Left from principal 2's mechanism. However if one of the two principals should deviate to any other mechanism, then the agents will instruct the non-deviator to minmax the deviator (which the non-deviator has committed himself to do when they make that recommendation). Since each of the agents expects the other two agents to make a unanimous recommendation, there is never an incentive for them to deviate in this continuation equilibrium.

This may be starting to sound somewhat abstract. However the important feature of this little example is to show that agents can reveal to one principal that another principal has deviated by sending quite simple messages. These messages allow one principal to respond immediately to the deviation in a way that makes it unprofitable. This is not an abstraction, it already happens in digital markets.

If you want to see what it looks like, here is a picture of a website offer at a website called Cheapair.com which offers low cost flights by aggregating offers on airlines' and other aggregators' websites:

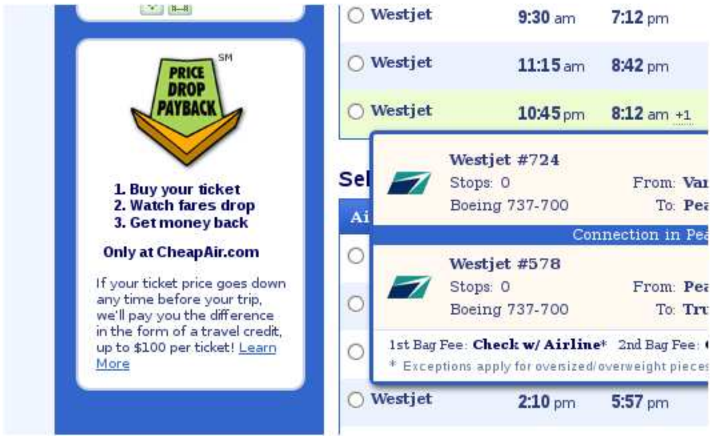

Notice the Cheapair Price Drop Payback - if one of Cheapair's competitors ever cuts price on a flight you book with them, they will refund the difference. This looks pretty much like something the IO literature used to call a *Meet the Competition* clause. It is obviously anti-competitive because it prevents cheapair's competitors from gaining market share by cutting price.

Characterization. Of course, this is just a simple example. It hides all the real world complexities that cheapair.com has to deal with. Yet it is suggestive of a couple of pretty important things. The Yamashita example given above is a *coordination game.* Coordination games typically have many equilibrium outcomes. This is true in the example, since the outcome (Down, Right) can be supported as an equilibrium just as easily as (Up, Left). This isn't simply a curious theoretical point. When mechanisms compete, as they do in digital markets, having lots of competitors and ease of entry will not ensure that equilibrium outcomes have any desirable welfare properties. As will become apparent below, large markets where a large number of sellers compete in mechanisms can support equilibrium outcomes that are very collusive. At least in the example above, the *First Welfare Theorem* (every competitive equilibrium is pareto optimal) does not hold, though an analog of the *Second Welfare Theorem* does seem right.

Second of all, the equilibrium in the game above suggests a familiar argument. If the principals designing mechanisms do the right thing, then they can end up with a pretty good payoff. If they deviate, they expect the actions of their competitors to change in a way that punishes them. This sounds much like the ancient minmax principle from complete information game theory. One of the advantages of the minmax approach is that minmax values can be found by solving straightforward maximization problems. If something like this minmax approach is true in more realistic problems with incomplete information, then it might be possible to describe the set of equilibrium outcomes using a set of inequalities instead of trying to write down a complex extensive form game and finding all its equilibrium outcomes. In fact, this is starting to sound a little like the good old revelation principle might actually be enough to describe the outcomes.

There are, however, two shortcomings in Yamashita's approach. The first is the less important, his theorem applies only to pure strategy equilibria in environments where players are either principals or agents, but not both, and principals never have private information. So his theorem doesn't really apply to digital markets where players will often play both roles.

The second shortcoming can be illustrated with a simple example due to Balazs Szentes ((Szentes 2010)). Again, there are two principals and 3 agents.[4] The payoff matrix is given by

1\2	Left	Right
Up	$1, -1$	$-1, 1$
Down	$-1, 1$	$1, -1$

The first payoff in each cell is the payoff for principal 1, the second for principal 2. The payoffs for each of the agents coincide with the payoffs for principal 1. So the agents want the outcome to be either (Up, Left) or (Down, Right). Principal 1 offers the Yamashita like contract that asks the agents to recommend an action and commits to carry out the action recommended by the majority. Suppose principal 2 offers a contract that specifies an action for each profile of recommendations made by the three agents. Then there is an equilibrium in which principal 2 simply commits to take action Left no matter what the agents say and the agents unanimously recommend Up to principal

[4]Szentes' example, like Yamashita's argument, requires three agents in order to support the majority report rules outcome. It should also be pointed out that this example exploits the restriction to pure strategies and non-random mechanisms that are part of Yamashita's environment. If the principals can commit themselves to mechanisms with random outcomes, this problem does not arise.

1. To see why this is an equilibrium, suppose that player 2 deviates to some other mechanism, and the agents send messages that support the action Right, for example. Then it will always be part of an equilibrium for the agents to switch their recommendation to principal 1 to Down. The interesting thing about this outcome is that principal 2 gets payoff -1 which is less than his minmax payoff.

What causes this is that agents see the deviation by principal 2 then recommend an action to principal 1 that minimizes principal 2's payoff when he takes that action. In other words, principals can only guarantee themselves their maximin payoffs. This isn't much of an issue in complete information game theory since the minmax and maxmin values for players are the same (provided punishments can be correlated). With incomplete information this presents an insurmountable complication for Yamashita's approach. The analog of the minimax in that case is to imagine that the principals are invited to join a mechanism in which a central planner will ask the principals to report types, then based on the information he acquires, he will instruct each player which action to take. If one of the principals refuses to participate, then the planner will force the other principal to carry out a punishment. A fundamental requirement for this to describe a set of incentive compatible and individually rational outcomes is that the punishment the planner imposes be independent of what the deviating principal who refuses the offer to participate subsequently chooses to do. If instead, the punishment is allowed to vary with the crime, we are essentially back to trying to describe an equilibrium of a particular extensive form game.

This raises a very basic question. Games in which principals compete by offering mechanisms are really just extensive form games of incomplete information. Normally we expect to be able to describe the equilibrium outcomes of such games by using the revelation principle. That is, an outcome function should be supportable as an equilibrium outcome in some competing mechanism game if and only if it is incentive compatible and individually rational. Individually rational means that there is a punishment that the others can impose on a player who refuses to participate that makes him worse off than he would be by participating whatever his payoff type, and whatever he decides to do after refusing to participate. Szentes' example illustrates that this kind of characterization may not always work.

To describe conditions under which such a characterization is possible, (Peters and Szentes 2012) study an abstract competing mechanism game in which contracts are finite strings of characters written in some language that includes all the usual arithmetic operations. Each principal's contract takes the strings written down by the other players as an input that is used to determine which action he will take. Formally, each player's contract is required to be a definable function of the other players' contracts.

The mathematical content of their theorem is to show that when contracts are required to be definable functions (there are more definable functions than there are Turing machines), then there must be a way for players to write contracts, so that no matter which of their pure actions they want to take, they can craft their contract (they might need to write different contracts for each of their pure actions) so that it elicits the same response from all their competitors. They call this the *invariant punishment property.* This seems at first glance to say very little. It doesn't say what this common response will be, only that there is some common response like this. However, it provides the key insight needed to characterize equilibrium outcomes in competing mechanism games.[5]

To see how, return to the Szentes example above in which there is an equilibrium in Yamashita style contracts in which Player 2 receives a payoff -1. If instead, the players compete by offering contracts that are definable functions of each other, then each player must be able to write a pair of contracts, one for each of his pure actions, that elicit the same response from his opponent. For Player 2, this response is going to be either Up or Down. Whichever it is, Player 2 can attain the payoff 1. So each player has a pure strategy min max value of 1. This proves that the game has no pure strategy equilibrium in definable contracts - just like the original game of matching pennies.

The important point in all this is that Yamashita like contracts will introduce new equilibria that cannot be understood using the minmax like property of the revelation principle because punishments implemented by participants on players who unilaterally refuse to participate can depend on what a non-participant eventually decides to do. What (Peters and Szentes 2012) shows is that competing mechanism games can be characterized using the revelation principle provided contracts are definable functions of each other.

The upshot of this is that very simple games can be constructed which will have all the same equilibrium outcomes as games in which players are allowed to use very complex definable functions as contracts.

[5]Definability is sufficient for invariant punishment but not necessary. Definability relies on action and type spaces being countable. Nonetheless, it is possible to support the invariant punishment property when the action space is uncountable by imposing arbitrary restrictions on feasible mechanisms.

Toward a Characterization and the Role of Institutions. I'll describe the main theorem in all this just to illustrate its generality and breadth. We begin with existing institutions that define property rights which essentially define the set of actions that are available to each of the players. We refer to these institutions as a default game. There are n players in this default game. Each player has a finite action set A_i and a finite set of possible payoff types given by 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. The notation Q is used to represent the set of all such mixtures. 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 $u_i : Q \times T \to \mathbb{R}$ where u_i is linear in q. An *outcome function* is a mapping $\omega : T \to Q$. So player i's payoff from this outcome function is $\mathbb{E}\left\{u_i\left(\omega\left(t\right), t\right)|t_i\right\}$.

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,

$$\mathbb{E}\left\{u_i\left(\omega\left(t\right), t\right)|t_i\right\} \geq \mathbb{E}\left\{u_i\left(\omega\left(t'_i, t_{-i}\right), t\right)|t_i\right\}. \tag{0.1}$$

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. A player who refuses to participate has to go back and pick an action in the default game. So the outside option is both type dependent and endogenous.

We allow the mechanism designer to implement a punishment that relies on information that has been collected from non-deviating players.

Let $\rho_i : T_{-i} \to Q_{-i}$ 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 $\{\rho_i\}_{i=1,n}$ such that for every player i,

$$\mathbb{E}\left\{u_i\left(\omega\left(t\right), t\right)|t_i\right\} \geq$$

$$\max_{a_i} \mathbb{E}\left\{u_i\left(a_i, \rho_{A_{-i}}\left(t_{i-i}\right), t\right)|t_i\right\}. \tag{0.2}$$

These are the incentive and individual rationality constraints associated with Bayesian equilibrium in Myerson's textbook (Myerson 1997). The main theorem in all of this literature is the following

Theorem 1. *((Peters 2010)) An outcome function ω is supportable as a perfect Bayesian equilibrium in a 'regular' competing mechanism game with seven or more players, if and only if it satisfies (0.1) and (0.2).*

The contribution of the theorem is twofold. First, it shows how to construct a single competing mechanism game that will support every incentive compatible and individually rational outcome as an equilibrium. This game is an abstract game, so I won't describe it here. However, it has two properties that are critical. First, mechanism designers only commit their own actions, and second, agents communicate with principals privately. No central planner is required to coordinate communication or instruct players about the actions they should take.

Second, notice that Myerson's characterization is of Bayesian equilibrium. Of course, any perfect Bayesian equilibrium in an extensive form game is also a Bayesian equilibrium. However, the if and only if requires the argument be made in both directions. An outcome that satisfies (0.1) and (0.2) requires a punishment by participating players. Punishments need not be sequentially rational in the sense once the participating players realize they are going to punish a deviator, they won't necessarily want to report their types correctly. This isn't an issue in a single principal problem since the principal can just ask agents to report their types before they realize there will be a punishment. Yet the competing mechanism by its very nature has agents observing deviations before they report. I'll return to this issue below to explain how the 'perfect Bayesian' qualifier gets into the theorem. For the moment, simply note that this makes the proof of this theorem much more involved than simply applying the revelation principle.

Notice that the theorem makes no distinction between principals and agents. Implicitly all players are able to make commitments about their own actions. It doesn't really make sense in this context to thnk of some players being principals and others agents. The distinction is based on knowledge of an institutional structure that is assumed to be missing here. For example, we might choose to model a market by assuming that one of the players owns an auction house and solicits bids and asks from the other buyers and sellers. We would call the owner of the auction house a principal and the others agents. This is all fine, except that what this paper is worried about is that some of the buyers and sellers might try to set up a competing exchange market. Even if we are absolutely sure that no such market exists, we still need to worry about what would happen if one of the 'players' we to deviate to set up an alternative auction house. If he did so, we would then have a second 'principal'. The view taken here is that the identity of principals and agents is determined endogenously - it isn't part of the description of an environment. For this reason we write the theorem using the term players.

Part of the implicit meaning of the term 'principal' is "... a player who receives messages from others but doesn't send messages...". Similarly, an agent sends messages but doesn't receive them. The remark made in the previous paragraph also applies here. Communication is part of the equilibrium, not part of the environment. So in the setup above it is assumed that all players can communicate (though they might choose not to in some equilibrium). If the players in the Theorem above are thought of as principals, then the assumption is that principals can communicate with each other, as well as with agents. So the composition of principals and agents among the 'seven or more' players is irrelevant.

Finally, note that the theorem is restricted to something called 'regular' competing mechanism games. This qualifier essentially restricts the theorem to games that satisfy the definability restriction in (Peters and Szentes 2012). This restriction is needed because Myerson's inequalities can not characterize the equilibria of all competing mechanism games.

I'll illustrate how to build a specialized game with an example below instead of using the general method.[6] As for the complexities that have plagued the existing literature, note that the types that are used in the characterization above are players' payoff types, not the more complicated types that embed all the players' market information.

One of the possible outcomes that can be supported in competing mechanisms is any Bayesian equilibrium of the default game. Notice that even if all the equilibria of this default game are desirable, many of the outcomes that satisfy the inequalities above are not going to be desirable.

The downside of the theorem is that it seems to say that it is going to be hard to predict the outcome in a competing mechanism game. Among the many outcomes that might be supportable are many that might have very bad efficiency properties. There can be no presumption

[6] The general method is described in (Peters 2010).

that competition will support good outcomes, or as we might say to undergrads, the First Welfare Theorem isn't going to hold. The failure of the first Welfare Theorem has nothing to do with externalities, it fails because simple minded notions of price competition just don't capture what is going on.

An example. As mentioned above, the competing mechanism game that is developed in the proof of Theorem 1 is an abstract game. However, to illustrate the implication of the theorem, it helps to go through an example that uses more familiar techniques.[7] This example also illustrates the sense in which institutions don't make much difference in all this. The example begins with an institution that is intended to promote efficiency. We'll illustrate how competing principals can work their way around this institution with appropriate mechanisms and communication in order to support a collusive outcome.

Of course the claim that collusion might occur in a large market is not new. The usual argument is informal, somehow sellers find a way to play cooperatively and implement some tacit agreement. The content of Theorem 1 in this context is to show formally how this tacit agreement is supported as a subgame perfect equilibrium using standard sorts of commitments.

This good institution in this example is just a double auction. The double auction is a sort of prototypical competitive market. The 'price' in the double auction is chosen to match demand and supply. Large double auctions always have Bayesian equilibrium which support outcomes that are ex post efficient, in the sense that the buyers and sellers who value the goods most highly will get them. We will demonstrate how sellers can undo this.

To do this, suppose that each of two sellers has a single unit of output to which he or she assigns a value of 0. There are two buyers, each of whom has a private valuation, either v_l or v_h ranked in the obvious way with $0 < v_l < v_h$. Each buyer is interested in acquiring a single unit of output. Payoffs to the seller are equal to the money he receives while payoffs to each buyer are equal to their private valuation when they succeed in trading, less the money they pay. We assume that valuations are correlated. To make life simple suppose that both valuations are the same with probability $\pi > \frac{1}{2}$ and that they are equally likely to (both be) v_h or v_l in that case. When the valuations are different (which occurs with probability $1 - \pi$), each of the two profiles occurs with equal probabability.

[7]This example is taken from (Peters 2010).

In the double auction that guides the interaction between them, players submit bids. The two available goods are awarded to the two highest bidders at a price equal to the third highest bid with the proviso that if there are more than two highest bidders, then the good is awarded to buyers whenever possible and randomly otherwise. For the purposes of illustration, focus on pure strategy equilibrium. There is a continuum of (ex post efficient) Bayesian equilibrium outcomes for this game in which all bidders bid $p \in (0, v_l)$ independent of type. The sellers' payoff in each of these equilibrium outcomes is p, high value buyers earn $v_h - p$ and low valuation buyers earn $v_l - p$. In all of these equilibrium outcomes trade occurs for sure. Evidently, the best outcome for sellers occurs when $p = v_l$. [8]

We are interested in characterizing the set of outcomes that could be supported in a competing mechanism game in which sellers offer their own contracts before the double auction offers. Taking the perspective of Theorem 1, we can first characterize the set of supportable outcomes. We can then take one of them and show how to construct a competing mechanism game to support it.

An outcome is a set of prices and trades that will occur for the each profile of types for the buyers. Let's leave the outcomes implicit, and do the characterization using the type contingent prices p_{hh}, p_{ll} and p_{hl} that prevail when both buyers have high values, both buyers have low values, or buyers have different values. To support higher prices, the sellers are going to have to find some way to restrict output. In particular, one of the sellers has to find a way to hold back his output when one of the buyers has a low type even though a profitable trade is possible.

The example also illustrates a way to use Theorem 1 in a more applied way. Instead of modeling a competing mechanism game that is supposed to describe the interaction between market participants, we can just write down the outcome that we are concerned about. In this context that will be the supportable outcome that maximizes the sellers' ex ante surplus. It is easy to find this since we don't have to do any game theory. We could then study its properties and compare them with the properties of the outcome function we want.

[8]There may also be inefficient asymmetric equilibria in which the trading price is above v_l with positive probability provided v_h is sufficiently large relative to v_l. In the rest of the discussion it will be assumed that v_h and v_l are close enough to each other that we can ignore these. I thank Gabor Virag for pointing out this issue to me.

Alternatively, once we understand the bad outcome as described by Theorem 1, we could then proceed to use some kind of structural estimation to recover the values of the buyers and sellers that support this bad outcome. A natural counter factual would then be to calculate the surplus we might have obtained with these same values had we been able to implement our most desired outcome - say the outcome that maximizes the ex ante surplus of buyers. This provides a natural way to calculate the costs of this collusive behavior.

So we proceed to describe the incentive constraints and the maximization problem.

A high value buyer doesn't know the other buyer's type. When he reports his type as high, he is accepting a trade at a price that can be either p_{hh} and p_{hl}. He should prefer this lottery to what he could get by pretending to be a low value buyer. This gives the completely standard incentive condition

$$\pi\left(v_h - p_{hh}\right) + (1-\pi)\left(v_h - p_{hl}\right) \geq$$

$$(1-\pi)\frac{1}{2}\left(v_h - p_{ll}\right). \tag{0.3}$$

He trades for sure if his value is high at a price that might depend on the value of the other buyer. Since one of the sellers is going to withhold his output when some buyer has a low type, he will fail to trade if he pretends to be low value and the other buyer has a high value. If the other other buyer's value is low, he will have the same chance to trade as the other buyer - $\frac{1}{2}$.

Similarly, the incentive condition for the low value buyer is

$$\frac{\pi}{2}\left(v_l - p_{ll}\right) \geq \pi\left(v_l - p_{hl}\right) + (1-\pi)\left(v_l - p_{hh}\right). \tag{0.4}$$

Since no one can be forced to trade, the individual rationality constraint is very simple here. It just requires that all buyers and sellers earn non-negative surplus.

Sellers' expected surplus assuming they are equally likely to be chosen to withhold output is given by

$$\pi\left(\frac{1}{2}p_{hh} + \frac{1}{4}p_{ll}\right) + (1-\pi)\frac{1}{2}p_{hl}. \tag{0.5}$$

The sellers' best expected surplus is found by maximizing this expression subject to the constraints above.

To understand the solution to this problem, observe that the low value bidder only trades at price p_{ll}. So he won't be willing to participate if $p_{ll} > v_l$. It is then immediate that $p_{ll} = v_l$ at the solution to this problem. Then most of the solution can be gleaned from the following

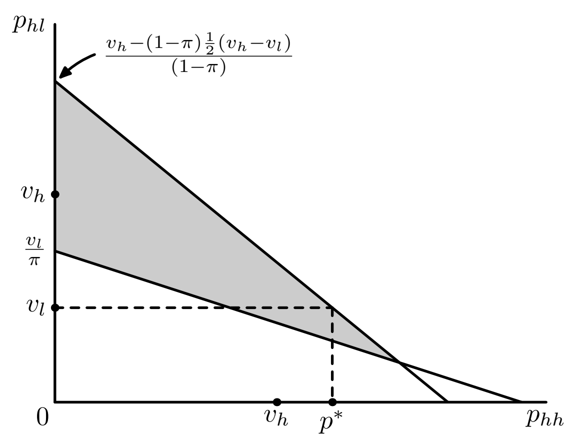

Figure 0.1. The Solution

Figure 0.1, which represents these two incentive conditions for the case when v_h and v_l aren't too different, and π is high.

The steeper of the two curves in figure describes the set of (p_{hh}, p_{hl}) pairs that make the high type buyer indifferent between revealing his type and pretending to be a low value buyer. This presumes that the price when both buyers claim to have low values is v_l. The high value buyer pays prices p_{hh} and p_{hl} when he trades. So if these are too high, he will be better off pretending to be low value and getting nothing. The set of price pairs that are incentive compatible for the high value buyer are those below the curve for this reason.

The flatter of the two curves[9] represents the set of price pairs that make the low value buyer indifferent between revealing his type truthfully and pretending to be high value. Reversing the reasoning above, if the prices (p_{hh}, p_{hl}) are too low, the low value buyer will want to pretend to have a high value so he can buy at these low prices. As a result, the prices that are incentive compatible for the low value buyer are those above this curve. The set of prices that are incentive compatible for both are those in the shaded area.[10]

The sellers' iso-profit function has the same slope as the steeper of the two curves, as is readily seen by comparing (0.3) and (0.5) above.

[9]The curves have different slopes because the low and high value buyers have different beliefs about whether or not the other buyer has a high value.

[10]This set gives all price pairs that are incentive compatible assuming that sellers hold back one unit of output when both buyers claim to have low values. All the price pairs on the line segment from $(0,0)$ to (v_l, v_l) are also incentive compatible and coincide with equilibria in the default game in which there is always trade. The sellers do better in the collusive equilibrium described here provided π is high enough.

As a result, any pair of prices on the upper right edge of the shaded triangle constitute a solution to the problem defined above.

To illustrate the kinds of arguments that are used to support this sort of collusion, focus on the outcome in which the price is p^* when both values are high, and v_l otherwise. This outcome is illustrated in the diagram. Since payoffs in this example are quasi-linear, prices don't matter when calculating overall surplus. The cost of this most collusive outcome is given by the expected value of lost trade, i.e., $\frac{v_l}{2}$.

The example continued - the competing mechanism game. To begin, suppose that Sellers 1 and 2 set up their own websites where they make offers before the double auction occurs. The website offers are commitments that are somewhat broader than a simple price offer. I'll use the technique I described above when I talked about the Cheapair.com website. The offer each seller makes is the same. They offer to sell at price p^*, but they also guarantee that if a buyer purchases the same good at a price that is no higher than p^*, then they will accept a return of the good for a full refund. Call this mechanism m^*.

There is also a third website in this story where both buyers and sellers submit comments at the start of the whole process. The comments are cheap talk, opinions, ratings whatever. For the purposes of this story, I'll imagine that the comments are similar to ratings, literally they will be numbers between 0 and 1.

The timing of the game will be that in the first stage, sellers will describe their offers. In the second stage buyers and sellers will post their ratings and buyers will decide whether or not the want to accept either of the sellers' offers. In the third stage buyers and sellers whose offers have not been accepted can bid in the double auction. To keep things short, it will be assumed that bidders in the double auction can see who they are bidding against.

Buyers make no commitments at all in this story, so it is a bit tedious to describe their strategies. They look like this:

- if both sellers offer the mechanism m^*, all four players should choose a number in $[0, 1]$ uniformly and report it to the feedback website. The fractional part of the sum[11] of all these reports is denoted by $\overline{x}$.
  - in any history in which all four players are participating in the auction, bid 0;
  - If type is low

[11]For example, the fractional part of 3.5 is 0.5.

    * bid v_l in an auction with 2 buyers and 1 seller, bid 0 if the only other bidders in the auction are sellers;
    * reject all offers, then stay out of the auction if one of the sellers' website offers was accepted.
    * If both sellers' offers were rejected, and you are buyer 1 stay out of the auction if $\underline{x} < \frac{1}{2}$, otherwise enter the auction and bid as above;
  - if both offers were rejected and you are buyer 2, stay out of the auction if $\underline{x} \geq \frac{1}{2}$, otherwise enter the auction and bid as above.
  - if type is high
    * accept one of the sellers' offers then bid p^* in the auction if the other buyer chooses to participate in the auction and v_l if the other buyer stays out of the auction.
- Finally if either seller offers a mechanism other than m^*, reject all offers, enter the auction and bid as above.

The seller's strategy goes like this

- in any history in which all four players bid in the auction, bid 0;
- bid v_l in the auction if the other seller's offer was accepted and only one buyer participates in the auction, bid p^* in this case if both buyers bid in the auction;
- if the other seller's offer was rejected, you are seller 1, and $\overline{x} < \frac{1}{2}$, give up and stay out of the auction. Seller 2 should do the same thing, except to replace $\overline{x} < \frac{1}{2}$ with $\overline{x} \geq \frac{1}{2}$.
- if the other seller's offer was rejected, you are seller 1, and $\overline{x} \geq \frac{1}{2}$, bid v_l in the auction if the other seller chooses not to participate and bid 0 otherwise. As above, use the same rule replacing substituting $\overline{x} < \frac{1}{2}$ for seller 2.
- if the other seller offers a mechanism other than m^*, bid as described above in the double auction.

The proof that these strategies constitute a perfect Bayesian equilibrium is somewhat tedious since there are a lot of deviations. The full proof is in (Peters 2010). Note here that they implement the sellers' most preferred outcome.

If both buyers have high types, they both accept the sellers' offers and both pay p^*. [12] If only one buyer has a high type, he accepts

[12] At this point, we'll just take it for granted that there is some implicit coordination device that ensures that each buyer accepts the offer of a different seller - for example, buyer 1 lives closer to seller 1, etc. Each seller is offering to buy at a

one of the sellers offers, but then buys the good a second time in the double auction (at price v_l). He then uses this sale to get a refund from the original seller. If both buyers have low types, they both reject the sellers offers. The sellers then use the numbers published on the ratings site as a public correlating device to determine which of the sellers will offer his good for sale in the double auction.

All the players send uniform noise to the ratings site. The public correlating device is created by adding up all the numbers then taking the fraction part of this sum. This simple device has a very useful property - the fractional part of the sum will be uniformly distributed on $[0,1]$ independent of the number that any particular player sends provided the others are all choosing their numbers uniformly on $[0,1]$.[13] As a consequence, it will always be a best reply for each of the players to choose uniformly no matter how the messages are being interpreted.

To illustrate the deviations, consider seller 1's options after a history in which both offers have been rejected. If the fractional part described above $\overline{x} < \frac{1}{2}$, he is supposed to stay out of the auction. Since $\overline{x}$ is public, he anticipates that seller 2 will participate in the auction along with one of the two buyers. If he chooses to participate, he expects that both the other participants will bid 0. So staying out of the auction is sequentially rational in this case.

Similar arguments cover the other cases - in particular the case where one seller deviates and offers some completely different mechanism in the first stage. This triggers a full participation auction in which everyone bids 0. This is an undesirable outcome for the deviating seller.

In the end, though the double auction has only efficient equilibria, the sellers can work around it to support another equilibrium.

Communication and Commitment. One of the advantages of this formulation is that the competing mechanism games provides a formal description of the two components of the strategic interaction between buyers and sellers that are common in all environments. Mechanisms are mappings from messages to whatever actions a player controls in a

price that is above their valuation, betting that they can get a refund and buy at a lower price if the other buyer has a low valuation. A buyer could offer to accept the offer of the seller who is supposed to sell to the other buyer, and thereby learn the other buyer's type. The payoff to that deviation is the same as the payoff to following the strategy of a low value buyer, which yeilds the high value buyer the same payoff as accepting the seller's initial offer. So the buyer cannot profitably deviate by offering to accept the offer of the wrong seller.

[13]The proof of this theorem can be found in the appendix of (Peters and Troncoso-Valverde 2013).

game. For them to work, players need to be able to make commitments and send messages.

One might ask whether there might be fewer potential outcomes if one were to simply remove these two components from the interaction. Removing commitment doesn't really make sense in a game theoretic setting. By definition a game describes a set of enforceable commitments. A player who chooses an action isn't allowed to change his mind and reverse it. When players participate in an extensive form game, they can commit themselves not to move out of turn.

All the existing literature on competing mechanisms essentially imposes restrictions on the set of messages that are verifiable. A message is verifiable if a principle can write a binding commitment based on that message. At one extreme, no messages are verifiable and principals simply play a non-cooperative game possibly influenced by cheap talk. In a common agency, messages are restricted to stating actions that might be taken by principals. The competing auctions literature makes bids verifiable. This paper basically assumes that all communications between principals and agents are verifiable.

A number of variants on these themes could be imagined. For example, it might seem sensible to try to restrict players to 'private' communication. It is far from obvious how to model private communication, but one way is just to imagine that players are linked together in a complete communication network in which there is an edge connecting every pair of players. A pair of players can transfer messages along this edge as often as they please.

Suppose there are four players and that player 1 wants to send the same message (a report about his type, for example) to each of the three other players, and that he wants to convince each of them that he has sent the same message to everyone. The following communication protocol will do the trick. Suppose that after player 2 receives player 1's message, he asks players 3 and 4 what messages player 1 sent them. He announces that if a majority of the messages that he receives agree, then he will believe that player 1 sent the same message to all three players and that this message was the same as the majority report of players 1, 3 and 4. Each of the players 2 and 3 believe that player 1 has sent the same message to all three players, and that the other player will report what he heard truthfully. Both 3 and 4 must then expect that player 2 will hear the same message from two other players. He might not like the message, but whether he does or not he can't do anything about it, so it will always be a weakly best reply for him to truthfully convey the message he heard from 1.

This device makes player 1's message look for all purposes like a public message even though all communication is private. (Peters and Troncoso-Valverde 2013) show how to use this to create correlating devices, and convey type information to other agents when communication is private. If players can make commitments based on whether or not enough players' messages agree, then this method can be used not only to correlate the actions of different players, but also how to implement randomized outcomes using non-random contracts.

A related question is whether one might not be able to restrict outcomes in competing mechanism games by assuming that some players cannot communicate at all. (Renou and Tomala 2012) study this problem when there is a single mechanism designer managing the actions of a group of agents. This mechanism designer wants to collect type information from all agents in order to do this, but is unable to communicate directly with some agents. (Renou and Tomala 2012) show that provided the principle is 2-connected to each of the agents, then anything that can be supported with full communication between the principle and agents can also be supported on the restricted communication network.

A network is 2-connected if there are two separate paths, possibly involving a number of the other agents, that any agent can use to communicate his type to the principal. If this is the case then an agent can encrypt his type, then send two separate messages, each containing the encrypted value of his type, the other a part of the encryption key needed to decode the type information. Any agent along the path who receives the message cannot understand it with only part of the encryption key. The principal, however, ends up with both messages, so is able to decrypt the type information. The same kind of result is likely true when are competing principals, though I don't know of any theorems on this.

Finally, the communication required to support some of these outcomes is not just extensive in the sense that it requires a rich network structure with free communication, it can also be intensive in the sense that some of the messages that players have to send are themselves complex. For example, the proof of Theorem 1 that appears in (Peters 2010) requires that agents recommend direct mechanisms that principals should use. It is hard to imagine agents being able to articulate such complex messages in practice. This, of course, is always an issue when using the revelation principle since buyers in markets usually have no idea what their 'type' is. In an interesting paper (Xiong 2013) explains how to support outcomes like those described above by having agents send messages that are elements of $[0,1]$. The idea is remarkably simple. Agents have to convey 2 bits of information to principals - the agent's payoff type, and his market information. Once the market has settled on a profile of direct mechanisms that principals are supposed to use, the only market information that is important from the principals' point of view is which of his competitors had deviated from this agreement. Agents can convey all this information by sending a element in $[0,1]$ if the principals interpret the information in a special way. First, the interval $[0,1]$ is divided up into j sub-intervals where j is the total number principals involved in the market. One of the subintervals is chosen as the default. If agents all send messages in this special sub-interval, then the messages are interpreted as meaning that all the principals offered the mechanisms they were supposed to offer. In that case, the principal commits himself to whatever direct mechanism he was supposed to use on the equilibrium path. Each subinterval represents one of the mechanism designer's opponents. If all the agents send messages in the i^{th} sub-interval, then the principal will interpret that to mean that principal i has deviated. The principal can then commit to implement the appropriate punishment against the deviating principal.

Each of these subintervals is, in turn, divided up into a set of subintervals, with each of these smaller subintervals representing the agent's possible payoff types. Since there are many ways to send the same type report, the actual messages can be used as randomized correlating devices provided the agents can communicate with each other (as in (Peters and Troncoso-Valverde 2013)). The effect of all this is that each agent can convey all needed information by sending a single number in the interval $[0,1]$ to each principle.

The upshot of all this is that if there are particular outcome functions that are especially undesirable (for example very collusive outcomes among sellers) then it isn't at all obvious what to do about them. The most obvious restrictions on mechanisms don't seem to have much effect.

Conclusion. Digital markets present something of a challenge. Traditional economic concepts and ideas seem ill-suited at best and misleading at worst when applied to digital markets. For example, there is nothing analogous to the first welfare theorem, though something akin to the second welfare theorem seems to apply, at least in the example above.

One of the most traditional approaches to problems like this is simply to respond that all the stuff that is supposed to go on as sellers promote these outcomes isn't really possible. Instead a very specialized model that is full of explicit and simple to understand restrictions on players' behavior will provide more insight than the "lots of equilibrium outcomes" approach in this paper. The double auction example described above was selected to dispel this idea. The double auction provides an explicit and simple description of the way players interact. The construction in the example was intended to show how sellers could work around this by using their own mechanisms. The complication in all this is that whether or not sellers are able to implement these workarounds is difficult to know. Whether they are or not isn't simply an issue about whether they are physically able to offer a meet the competition clause. There are many ways that communication mechanisms can accomplish the same thing without explicitly 'meeting the competition'.

Consider, for example, a very simple question: A seller might want to know whether a buyer has checked his competitor's website for a lower price? Recall that the basic argument in this paper is that sellers might want each other to do this in order to reduce the incentives for anyone to lower price. It might seem that a seller can't ask this question because the internet cookies that websites use to track consumers' search behavior on the internet can't be passed between different domains (at least given the way the current html spec works). There is a simple way around this. Each of the websites typically contains a little javascript button which allows you to 'like' a page on Facebook, or to 'Follow' a website on Twitter. When one of these buttons is loaded, it typically calls a script that it contained on the Facebook or twitter website. As this script is being loaded, the browser will accept a cookie from Twitter or Facebook that can contain the url and page loaded on the site along with the time the page was loaded. Then, when the buyer loads another page on a different website, a script on that website can again try to render the Facebook Like button. As the script is returned Facebook can read its own prior cookie so that it knows that the buyer has visited both websites. By selling back this information to either of the sites involved, it can trigger a meet the competition response.

To understand whether or not this is going on requires a much deeper knowledge of network topology than traditional economic models provide. It would be nice to know whether network structure provides any insight at all into the ways these markets operate. For example, as described above, the set of outcome functions a single principal can implement on a 2-connected network is the set of outcome functions that can be implemented on a fully connected network. It would be nice to know whether the same kind of result is true on a network of competing mechanisms. In particular 2-connected is sufficient for full implementation. It would be nice to know whether something like it is also necessary.

When discussing this problem, it seems natural to assume that it is sellers who are using digital tricks to price discriminate. This is partly because people have most of their experience the sellers' websites, and partly because economics has a tradition of assuming that sellers' make offers to buyers. Yet buyers can play the game as well. For example, the Facebook - Twitter trick described above can be circumvented by buyers if they use aggregator sites (for example Priceline.com or travelocity) to get price quotes instead of visiting independent websites.[14] Of course, as we explained above with the Cheapair.com example, this can also work to sellers' advantage. Yet it does suggest that more attention should be paid to digital techniques that lower prices.

One of the arguments that appears in pretty much every paper on competing mechanisms is the 'tipping' idea - buyers and sellers all want to be in the same market because that maximizes the potential gains to trade. Once they all get to one market, they get stuck because no one wants to go to a market where no one goes. In a digital market this idea makes little sense. Digital robots can scan all markets and instantly respond when an entrant makes an offer in that market, even if everyone expected that market to be empty. This begs the question whether there are methods that can be used to train markets to play the right equilibrium.

References

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

Martimort, D., and L. Stole (2002): "The Revelation and Taxation Principles in Common Agency Games," *Econometrica*, 70(4), 1659–1673.

McAfee, P. (1993): "Mechanism Design by Competing Sellers," *Econometrica*, 61(6), 1281–1312.

Myerson, R. (1997): *Game Theory The Analysis of Conflict*. Harvard University Press.

Pavan, A., and G. Calzolari (2009): "Sequential Contracting with Multiple Principals," *Journal of Economic Theory*, 50, 503–531.

Peters, M. (2001): "Common Agency and the Revelation Principle," *Econometrica*, 69(5), 1349–1372.

——— (2010): "On the Revelation Principle and Reciprocal Mechanisms in Competing Mechanism Games," Discussion paper, University of British Columbia.

[14]There are also firefox extensions that will turn off twitter and facebook cookies. Though these are useful, using them is a bit like playing 'whack a mole'.

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

Peters, M., and C. Troncoso-Valverde (2013): "A folk theorem for competing mechanisms," *Journal of Economic Theory*, 148(3), 953 – 973.

Renou, L., and T. Tomala (2012): "Mechanism Design and Communication Networks," *Theoretical Economics*, 7, 489–533.

Szentes, B. (2010): "A note on Mechanism games with multiple principals and three or more agents by T. Yamashita," Discussion paper, London School of Economics.

Xiong, S. (2013): "A Folk Theorem for Contract Games with Multiple Principals and Agents," Discussion paper, Rice Univeristy Discussion Paper.

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