CAN MECHANISM DESIGNERS EXPLOIT BUYERS' MARKET INFORMATION

MICHAEL PETERS

ABSTRACT. Competing mechanism games can have many equilibrium outcomes when mechanism designers can use mechanisms which ask agents to report deviations. Existing folk theorems assume agents have perfect information about designers' mechanisms, that they can costlessly convey very complex messages, and that they use weakly dominated strategies. This paper addresses two questions. The first is whether the basic logic behind the folk theorems for competing mechanisms can work without these assumptions. The second is to ask how large frictions have to be before collusive equilibria break down.

1. INTRODUCTION

The website a seller uses in an on line market represents a nice example of a *mechanism* as it is understood in textbooks on mechanism design. The website processes messages from buyers and these messages ultimately determine the trading price that a buyer is offered. The website acts as a pretty much perfect commitment - it is just a computer program that responds mechanically to messages and allows almost no ex post flexibility. In the tradition of a direct mechanism, the price that is offered to any buyer could well depend on messages that are received from other buyers. Indeed standard website driven pricing schemes, like airline ticket pricing, already incorporate this feature. All this suggests that a useful way to look at online markets is to view them as *competing mechanism games.*

The literature on these games has an ambivalent message. On one hand, the literature on competing auctions suggests that large markets will work efficiently and that mechanisms will be as simple as they can be.[1] For markets, the messages is the usual one, that competition is good.

June 16, 2016.

[1] See (McAfee 1993), (Peters and Severinov 1997), (Peters 2001b), (Virag 2010)

Alternatively, by exploiting the possibility that mechanisms can condition on one another, there is a literature showing that almost anything can happen in very competitive markets if mechanism designers use rich enough message spaces.[2] In particular, very competitive electronic markets can support very inefficient equilibrium outcomes.[3] This is really a surprising comment on one of the main lessons of game theory and mechanism design - it is advantageous to commit. Markets seem to work better when sellers don't (or can't) take advantage of all their commitment abilities.

The difference between these two literatures is an assumption about message spaces. In the former, buyers can freely convey information to sellers about their willingness to pay (their payoff types) by bidding, but are assumed not to be able to convey any other information. In the latter, buyers can costlessly report any information they have. In particular, buyers can see the mechanisms used by all sellers (as they can in the literature on competing auctions) and reveal all this information to all the sellers at the time they report their payoff types. Sellers can then design punishments whenever buyers report deviations from some collusive outcome.

Neither of these extreme assumptions is particularly satisfactory. If it is so easy for buyers in the competing auction literature to see all sellers' mechanisms then accurately report their payoff types, why is impossible for them to report anything else? Indeed, in modern digital markets buyers do convey more information. Sellers keep detailed records of buyers' search behavior. Since this behavior is related to information buyers receive as they search, it conveys market information in exactly the method assumed in the literature on the folk theorem.

On the other hand, it isn't hard to find cracks in the assumptions that the competing mechanisms literature uses to support the folk theorem. The first is that many consumers just don't know or understand how sellers' mechanisms work.

Here is a quote from the Wall Street Journal in 2011[4] describing airline ticket pricing.

> Though prices fluctuate frequently and the ups and downs
> of airline prices can frustrate and anger consumers, air-
> line pricing actually does follow a cycle during the week.

---

[2]The many papers include (Epstein and Peters 1999), (Peck 1994), (Martimort and Stole 1998), (Peters 2001a),(Xiong 2013),(Watson 2012),(Peters and Troncoso-Valverde 2013), (Peters and Szentes 2012),(Forges 2012) or (Peters 2015).

[3]This represents a kind of failure of the first welfare theorem.

[4]"Whatever You Do, Don't Buy an Airline Ticket On ..." by Scott McCartney, Wall Street Journal January 29, 2011.

> Many sales, in which some seats are discounted by 15%
> to 25% typically, are launched Monday night....

What is odd about this quote is that it articulates the fact that consumers are surprised by airline pricing and that the only way to understand it is to look retrospectively at the way prices behave. Evidently, airlines are in no hurry to explain their pricing strategy to consumers - which seems a violation of the idea that it is advantageous to commit (and let people know you have committed).

In the terminology of game theory, buyers in on line markets like the airline market have imperfect information in the sense that some of them don't observe the previous moves of mechanism designers - i.e, they don't understand the computer programs mechanism designers are using to set prices. If buyers can't see or understand deviations, there is no way they can report them to a mechanism designer. Buyers' market information seems quite different from their payoff information, which they surely do see and understand.

The second issue is that even if buyers do see and understand the mechanism a seller is using, it can be quite difficult for them to convey that information to another seller because of the built in infinite regress that is involved in describing a mechanism.[5]

The point of this paper is to consider a middle ground. In the model we describe below, information is both incomplete and imperfect. Sellers aren't sure whether buyers understand their mechanisms in the first place. So they aren't sure whether the incentive effects they have built in will be effective. We'll basically assume that buyers convey their payoff information, and their market information at different times.

We do allow the possibility that the cost of conveying these two types of information is different, but not because market information is complex and requires a complex message. We'll assume that the way buyers convey their market information is by using a technology that is known to exist - an html cookie. We'll assume the cookie is just a simple binary message that tells a seller whether a buyer who visits their website has done so before. Clearly, such a message involves no difficulties with infinite regress.

Even though it is very easy to leave cookies on websites they visit, we won't assume that it is free. This market message will involve an incremental cost. We'll use these costs to parametrize the frictions, small as they might be, in a digital market. Most of the logic in the

---

[5]The fact that can be hard to describe is whether a mechanism depends on whether another mechanism depends on whether the first mechanism depends ..... . See (Epstein and Peters 1999) or (Peters and Szentes 2012).

paper is centered around trying to figure out how large these costs have to be before a market full of inefficient outcomes supported with complex contracts collapses to a market whose equilibrium resemble the efficient outcomes of the competing auctions literature.

This issue is of some interest, as buyers convey all their information in the same way, by visiting sellers' websites. As we will show, restricting buyers to simple messages and assuming that buyers have imperfect information about mechanisms isn't restrictive at all - if buyers can send these messages costlessly, the market will support many equilibrium outcomes and the conclusion of the competing mechanism literature prevail - there are many equilibrium outcomes, some quite inefficient.

However, message costs do play a role. Obviously, if the costs of sending messages is high enough, buyers won't find it worthwhile to convey market information. By itself, that is unsatisfactory because it ignores the fact that high costs will also make it impossible for buyers to convey their payoff type information.

Our main contribution below is a theorem that shows how the set of supportable equilibrium outcomes depends on communication costs. As these costs increase the set of supportable prices shrinks and eventually disappears. The set of equilibria with prices above the 'competitive' price disappears entirely when a certain component of communication costs reaches approximately $\frac{1}{3}$ of the value generated by trade with the highest value buyer in the support of the distribution of buyer values. This seems to mean that communication costs must be very high before collusive equilibria disappear.

1.1. **Environment.** Since the main message of the paper is about communication the basic economic properties of the model are very simple. There are n buyers and m sellers with $m > n$. Each seller has a single unit of a homogeneous good to sell. Each buyer wants to acquire exactly one unit. The sellers' goods are all perfect substitutes. The value to a buyer v_i who trades at price p is $(v_i - p)$. Buyer values are identically and independently distributed according to some monotonic atomless distribution F such that 0 is in the support of F.

We'll assume that sellers all have the same value (i.e. cost) $c > 0$ and that this fact is common knowledge. Sellers who trade at price p earn surplus $p - c$. This is an extreme assumption, but fits with the idea that the good being traded is a well known commodity rather so that the buyers who enter the market are simply looking for trading partners rather than information about what the price might be. We'll briefly

discuss how changing this assumption might affect outcomes later in the paper.

The market is broken into three parts. This first is a discovery phase in which buyers learn about the various options available in the market. Presumably, buyers learn about potential trading partners by visiting their websites. We'll assume that when they do, sellers leave an html cookie on the buyers' browsers which they can recover any time the buyer revisits their website. We aren't going to model this procedure explicitly, we simply assume at the end there are n buyers who go through this process and end up being the demand side of the market.

The second phase is an exploratory phase where buyers can continue to browse sellers' websites. This is the phase where buyers are going to convey any information about the market that they received in the discovery phase, so this is where most of the analysis in the paper is focused.

This exploratory phase involves costly search for buyers. Buyers know (or suspect) that sellers use mechanisms that make their subsequent offers depend on this exploratory phase. Their search decisions are assumed to reflect a best reply to a belief about these mechanisms that is correct in equilibrium. So there is nothing behavioral in the story here.

Finally, we'll assume there is a trading phase in which buyers reveal their payoff types. Presumably buyers continue to visit websites looking for offers to trade. This could be a complex and potentially frictional process. However, since we are focused on market information, we'll black box this process and assume that buyers and sellers engage in a sellers' offer double auction which determines all the trades.[6]

The way the sellers' offer auction works is that each buyer submits a bid and each seller submits an ask to a centralized mechanism which then arranges all trades. The buyers and sellers with the m highest bids and asks end up with a unit of the good at the end of the auction, with ties going in favor of buyers. Each buyer who buys and each seller who sells pays (or receives if they are a seller) a price which is equal to the $m + 1^{st}$ highest of the bids and asks.

As $m > n$, there is always a non-trivial equilibrium in this auction in which buyers bid their values, and sellers ask c. The reasoning is two

---

[6]Effectively, we are making the assumption that is common to all of mechanism design which is that buyers can costlessly convey information about the values. We could make it costly to convey this value information as well. This would cause overall trade to be less efficient since the highest value buyers might not always be served. Since this sort of analysis is already well understood there seems little point in analyzing it in detail.

fold. Since the $m + 1^{st}$ bid or ask determines the trading price, and the highest m bids or asks always end up with a unit of the good, the only way the buyer can affect the trading price is by submitting a bid which does not allow him to trade. By the usual argument, bidding value is a weakly dominant strategy for buyers.

For sellers, if the seller expects the other sellers to ask c, then asking more than c will ensure that the seller is always one of the m highest bidders, because $m > n$.

This equilibrium is efficient because it ensures that all buyers who value is above c end up trading.

In the terminology of *(Peters 2015),* we'll think of the double auction as the *default game* and assume that the competing mechanism game is built on top of it. We'll adopt the restriction that only sellers can create mechanisms, so the sellers are the principals and the buyers are the agents in the competing mechanism game.

Sellers' mechanisms are programs that determine the ask prices the sellers will submit to the double auction. These programs process messages received from buyers during the exploratory phase of the market. As mentioned, buyers visit sellers during the discovery phase, at which point the seller is allowed to place a cookie on the buyer's browser. We'll assume that this cookie just records the fact that the buyer has visited.

If a buyer returns to the seller's website during the exploratory stage, then the seller knows the buyer has visited before. Some buyers will revisit, some won't. So we'll assume that the messages the seller uses are just binary. The message 1 means the buyer chose to revisit the sellers' website during the exploratory phase. The message 0 means that he didn't.

We'll assume that before everything else, the seller writes a computer program that records that number of buyers who revisited his website during the discovery phase, then submits an ask price to the double auction that depends on this number.

Buyer website visits during the discovery stage are costly. But from the perspective of the analysis in this paper, these costs are sunk by the time the exploratory phase occurs. The set of buyers who choose to bear this cost is exactly the set of n buyers we have assumed will participate in the double auction.

In the discovery phase, we'll assume that buyers discover the whole market - in other words they visit every seller's website. We'll extend this idea to the exploratory phase, and assume that there is a fixed cost to re-visiting websites. Once a buyer chooses to bear this cost, he can (and will want to) revisit all sellers' websites.

Finally, the double auction requires additional interaction between buyers and sellers which may be costly. We don't model this explicitly. Instead we just assume that all the buyers in the market choose to bear this cost and participate in the double auction.

If t is the number of buyers who revisit seller j during the exploration phase, then $a_j(t)$ is seller j's ask price in the double auction. This is all done programmatically. Sellers are perfectly committed and never make any decisions beyond the program that they write at the beginning of the process. In a terminology common in mechanism design, the message space for sellers is just $\{0,1\}$- 1 means visited, 0 means never returned. This message space is about as simple as it could possibly be. This is in contrast to the assumptions in most of the competing mechanism literature.

The second part of the model is intended to restrict what buyers understand about sellers' mechanisms after the discovery phase. We'll assume that what buyers see when they make their initial visit to any seller's website is the seller's opening offer $a_j(0)$. Buyers realize that sellers are using mechanisms to determine their ask prices, but as in the airline example in the first section of the paper, none of the buyers will know for sure what these mechanisms are, though they will correctly guess these mechanisms in equilibrium.

Finally, we'll further restrict buyers ability to observe sellers' mechanisms. We'll assume that the probability that a randomly selected buyer sees (or notices, or remembers) the initial offer at any website he visits is given by $\gamma(n)$, with $\gamma(n)$ being non-increasing in n. The assumption that only certain buyers notice this offer is in the spirit of the literature on inattention[7], particularly random inattention.
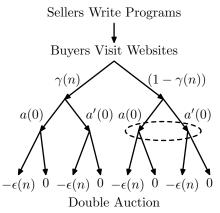
Sellers decisions are pre-programmed. However, at the point where the seller's program submits a bid to the double auction, it has incomplete information about whether any buyer knows what they are doing.

At this point, it could be so unlikely that buyers learn anything about price quotes, that there is no market information to convey. At the extreme, if buyers and sellers believe that no buyer observes anything about the seller's mechanism, then there are no commitments, and the equilibrium will just be the equilibrium of the double auction. This equilibrium will be ex post efficient when there are enough buyers and sellers. So we will restrict attention in what follows to environments where there is some minimal amount of market information for buyers to convey. In particular we are going to assume that for all n,

---

[7]For example, (Masatloglu 2015) or (Masatloglu, Nakajima, and Ozbay 2012)

$\gamma(n)(n-1) \geq 1$. In words, we are going to assume that on average, each seller believes that there is at least one buyer who sees each seller's initial offer. We emphasize that the seller has no idea who this buyer is.

This figure summarizes the basic properties of the model. After sellers write their programs, buyers make their visits to websites. $\gamma(n)$ is the probability with which any buyer notices the sellers initial price offer $a(0)$. With complementary probability $1 - \gamma(n)$ the buyer won't be able to detect a modification from $a(0)$ to $a'(0)$. Buyers then have to decide whether or not to bear the cost $\epsilon(n)$ to revisit sellers' websites. To end it all, the default game - a double auction - is played, with the sellers' programs submitting their ask prices.

One obvious equilibrium of this competing mechanism game occurs when sellers ignore buyer messages entirely and all players simply play the double auction non-cooperatively.

The issue in this paper is whether there might be equilibrium outcomes in which sellers set prices above c. If p is such a price, and all firms ask p in the double auction, then each firm has expected payoff

$$(p-c)\sum_{t=0}^{n} \frac{n!}{t!\,(n-t)!} \left(1 - F\left(p\right)\right)^{t} F\left(p\right)^{n-t} \frac{t}{m} =$$

$$(p-c)\frac{\left(1 - F\left(p\right)\right) n}{m}.$$

Of course, this isn't an equilibrium in the double auction as any firm can cut price slightly and sell for sure. The essential property of competing mechanism games is that when one firm cuts price other firms will modify their own prices. Our first objective is to check whether this logic survives when buyer have imperfect information about sellers' mechanisms, and find it costly to convey information.

1.2. **How high prices are supported as perfect Bayesian equilibrium.** We'll start by showing how sellers can use buyers' market information to maintain a high price equilibrium. The argument is a variation of the 'meet the competition' argument ((Salop 1986)). Notice a couple of things about it. First, the argument is fully symmetric - all firms use contracts. Yet none of the complex problems associated with infinite regress[8] emerge.

Second, note that the argument mimics the logic of a repeated game. The argument that this logic is sufficient to characterize all the equilibria in a competing mechanism game is due to (Peters and Szentes 2012).

Finally note that the messages (and for that matter contracts) that implement this equilibrium are all trivial. Contrast this with (Yamashita 2010), (Epstein and Peters 1999) or (Peters and Szentes 2012, Peters and Troncoso-Valverde 2013).

Let p be some price larger than c. Suppose that all sellers use the program or rule

$$a(t) =$$

(1.1)
$$\begin{cases} p & t < 2 \\ 0 & \text{otherwise} \end{cases}$$

to set their ask prices, where t is the number of buyers who re-visit their website.

This rule triggers price cuts after only two buyers choose to revisit. We'll consider other kinds of price trigger rules below, however, this one will play a role in all that follows. To distinguish this rule from others, we'll refer to it as a reciprocal pricing rule.[9]

We'll use the 'punishment' price 0 here because the point is to look for conditions under which these sorts of equilibria can't arise. A punishment price 0 gives buyers the biggest incentive to report deviations, so should support the biggest set of equilibria. If this seems implausible, the arguments below are readily modified to use, for example, c as the punishment price. The idea behind the reciprocal pricing program is that each seller will advertise a price p and adjust it only if 2 or more buyers re-visit his website. Buyers search all websites initially

---

[8](Epstein and Peters 1999)

[9]This term is also used in (Peters 2015) in a model in which firms can directly condition their actions on other firms contracts. The contracts here are a kind of special case in which the dependence comes indirectly through buyer messages. So we'll continue to slightly misuse the term here.

just because they need to know the set of sellers from whom they can purchase. However, when they see a deviation, they may decide to revisit websites in the hopes that prices will be driven down.

Proposition 1. *Let p satisfy $\frac{n}{m}\left(1 - F\left(p\right)\right) \geq \frac{1}{2}$, and suppose there is a v such that*

$$(1.2) \qquad \left(\frac{n-2}{n-1}\right)^{n-2} > \max\left[\frac{\epsilon\left(n\right)}{p}, \frac{\epsilon\left(n\right)}{v}\right].$$

Then there is a perfect Bayesian equilibrium in which all trades occur at price p and all sellers use the contract described by (1.1).

Proof. Let $v\left(n\right)$ be the minimum value of v such that

$$\left(n-1\right)\left(\left(1 - F\left(v\right)\right)\gamma\left(n\right)\right)\left(1 - \left(\left(1 - F\left(v\right)\right)\gamma\left(n\right)\right)\right)^{n-2} \geq$$

$$\max\left[\frac{\epsilon\left(n\right)}{p}, \frac{\epsilon\left(n\right)}{v}\right].$$

By (1.2), the set of v satisfying the inequality above is non-empty. Since both sides of the inequality are continuous in v, the minimum value exists. If $v\left(n\right) > 0$ then

$$\left(n-1\right)\left(\left(1 - F\left(v\left(n\right)\right)\right)\gamma\left(n\right)\right)\left(1 - \left(\left(1 - F\left(v\left(n\right)\right)\right)\gamma\left(n\right)\right)\right)^{n-2}$$

$$(1.3) \qquad = \max\left[\frac{\epsilon\left(n\right)}{p}, \frac{\epsilon\left(n\right)}{v\left(n\right)}\right].$$

The binomial probability term on the left hand side of the equation above is then the probability of exactly one success in $n - 1$ trials, when the success probability is $\left(1 - F\left(v\left(n\right)\right)\right)\gamma\left(n\right)$. This probability attains its maximum when 1 is the modal value of the binomial distribution. This occurs when $\left(1 - F\left(v\right)\right)\gamma\left(n\right) = \frac{1}{n-1}$, at which point, the probability attains value $\left(\frac{n-2}{n-1}\right)^{n-2}$. Then if there is a strictly positive v in $V\left(n\right)$, it must be that $\left(1 - F\left(v\left(n\right)\right)\right)\gamma\left(n\right) \geq \frac{1}{n-1}$ since by the properties of the binomial probability term

$$\left(n-1\right)\left(\left(1 - F\left(v\right)\right)\gamma\left(n\right)\right)\left(1 - \left(\left(1 - F\left(v\right)\right)\gamma\left(n\right)\right)\right)^{n-2} \leq$$

$$\left(n-1\right)\left(\frac{1}{n-1}\right)\left(1 - \frac{1}{n-1}\right)^{n-2}$$

whenever $\left(1 - F\left(v\right)\right)\gamma\left(n\right) > \frac{1}{n-1}$.

The strategy we want buyers to use is bid their values in the double auction without re-visiting sellers' websites if they do not observe any opening offers, if every opening offer they see is equal to p, or if their value is less than $v\left(n\right)$. Otherwise, they should revisit sellers' websites then bid their values in the double auction. As buyers who trade cannot

affect the trading price in a seller's offer double auction, bidding value is a weakly dominant strategy.

If a buyer finds himself in an information set in which he has not observed any opening offers different from p, then he should anticipate that if he does visit websites, he will be the only one to do so. The reason is that all sellers are expected to offer p, and so no other buyers will re-visit websites. As the buyer believes the seller is using a program like the one described by (1.1), he believe that if he unilaterally revisits, he won't affect prices. So it is a best reply for each buyer type to refrain from revisiting in any information set in which he does not observe an opening offer that differs from p.

Finally, if a buyer does see a seller post an opening offer different from p, then he expects that other buyers will re-visit websites if they have also seen the deviation, and if their values are larger than $v(n)$. If two or more of the other buyers choose to revisit, prices will fall whether the buyer visits or not. If no other buyers choose to revisit, the buyer will also have no impact on price. The only event in the seller can impact price is if exactly one of the other buyers chooses to revisit websites. This occurs with probability given by the left hand side of (1.3).

If $v(n)$ is smaller than p, then the buyer of type $v(n)$ would not trade on the equilibrium path. Since the punishment pushes the price to 0, his gain is $v(n)$ when he triggers the punishment. So he is just indifferent ex ante. Buyers whose values are higher than $v(n)$ will strictly prefer to try to trigger punishment by revisiting, while those whose values are less won't.

If p is smaller than $v(n)$, then the buyer of type $v(n)$ will trade in equilibrium. If his re-visit triggers the punishment he will gain p as a result. In this case, all buyers whose values are at least p will be indifferent about revisiting, while those whose values are below $v(n)$ will not be interested in re-visiting.

Finally, for a seller to an opening offer slightly below p will ensure that the deviating buyer sells at price p for sure if the deviation remains unreported. The probability that the deviation remains unreported is
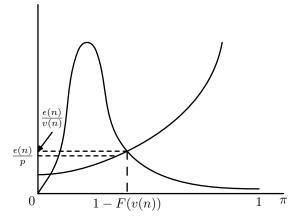(1.4)
$$n(1 - F(v(n))) \gamma(n)(1 - (1 - F(v(n))) \gamma(n))^{n-1} + (1 - (1 - F(v(n))) \gamma(n))^{n}.$$

This is just the probability that either no buyers choose to re-visit, or only one does.

As argued above, $(1 - F(v(n))) \gamma(n) > \frac{1}{n-1} > \frac{1}{n}$, so the modal value of the binomial distribution faced by firms is at least 1. As a

consequence (1.4) can be no larger than $\frac{1}{2}$. Since $\frac{n}{m}(1 - F(p))$ is at least $\frac{1}{2}$, deviations for firms are unprofitable.   □

The following diagram may help to explain some of the conditions in the theorem.

The single peaked function gives the probability

$$n\pi\gamma(n)(1 - \pi\gamma(n))^{n-2}$$

for values of π between 0 and 1. For any π, the corresponding buyer type $\hat{v}(\pi)$ that induces that π is just the solution to $\pi = (1 - F(v))$. So the higher π is, the lower is the corresponding value $\hat{v}(\pi)$.

The upward sloping function is $\frac{\epsilon(n)}{\hat{v}(\pi)}$. Then the solution to (1.3) coincides with the highest π (or lowest v) at which the corresponding buyer type is just indifferent between re-visiting website and not.

Notice that if the cost of re-visiting websites is 0, then all buyers will report, and the no deviation condition becomes

$$n\gamma(n)(1 - \gamma(n))^{n-1} + (1 - \gamma(n))^n \leq \frac{n}{m}(1 - F(p))$$

Under our assumptions, the expect number of buyers who see a deviation $\gamma(n)n$ is at least 1, the median of the distribution of the number of buyers who see the deviation is at least 1. The two terms on the left hand side of the expression about can therefore be no larger than $\frac{1}{2}$. So if it is costless to report, there will be an equilibrium supporting price p provided $\frac{1}{2} \leq \frac{m}{n}(1 - F(p))$, the condition in the theorem. One thing this illustrates is that even if messages are costly, high price equilibrium cannot be support if p is too high or if γ is too low. Of course, if the condition fails when it is costless to send messages, it also fails when messages are costly. So for the rest of the paper, we'll restrict attention to p such that $1 - F(p) \geq \frac{m}{2n}$.

Though this theorem illustrates how easy it is to support high priced equilibria using buyers' market information, we are more interested in what happens when costs are higher.

Some Game Theory. As described above, the competing mechanism game has imperfect information in the sense that most buyers do not directly observe the mechanisms that sellers are using, even though, in the usual game theoretic sense, they will correctly guess the mechanism in equilibrium. Even those buyers who do observe something about mechanisms don't learn very much - they only learn the opening offer. They believe that sellers adjust their ask prices in response to visits.

When the seller chooses the program used to set an ask price, he is faced with a mechanism design problem. However, it differs from the standard mechanism design problem in two ways. First, to decide on the best program to use, the seller needs to evaluate the alternatives. These alternatives are different mechanisms. When the seller 'deviates' to one of these, the deviation has the usual incentive effects, but only for those buyers who see the deviation. So two buyers with the same value will respond to these deviations differently depending on whether or not they see them.

Secondly, some deviations are never going to be observed by any buyers. For example, instead of cutting price when two buyers revisit, the seller could evaluate an alternative in which he cuts price after three buyers visit. If he tries this, none of the buyers will actually see it, so their behavior won't be modified.

Though the competing mechanism game we describe here has lots of equilibrium, these special properties impose a lot of structure on these equilibria.

Theorem 2. *In every symmetric equilibrium of the competing mechanism game described above, buyers re-visit websites with probability 0.*

Proof. Suppose to the contrary there is an equilibrium in which sellers use some common pricing rule $a(\cdot)$, and buyers re-visit sellers' websites with positive probability. Since every buyer who revisits, revisits all websites, all sellers set the same ask price $a(t)$ if $t > 0$ buyers revisit. Since no buyer sees $a(t)$, any modification to the seller's program that alters this price will have no impact on buyers' decisions about whether or not to revisit and, in particular, no impact on the probability that exactly t buyers revisit.

Then if $a(t) < c$, the seller obviously wants to deviate and raise his ask price above c, because he doesn't want to trade in that event. If

$a(t)$ is greater than c, then the seller's expected payoff when t buyers revisit is $(a(t) - c)(1 - F(a(t)))\frac{n}{m}$. The seller then has an incentive to lower his ask price in that event slightly below $a(t)$ so that he trades for sure. So sellers will deviate and rewrite programs whenever $a(t)$ is different from c.

So if there is an equilibrium where some buyers re-visit websites, it must be that $a(t) = c \forall c > 0$. In that case it can't be optimal for buyers to revisit websites because their visits won't affect prices.   □

This result says that website re-visits only play a role off the equilibrium path. If an equilibrium with $a(0) > c$ can be supported at all, then it must be because a continuation equilibrium following a deviation in which buyers re-visit websites, triggering price cuts by sellers. Furthermore, because there are no website re-visits in equilibrium, the only way a seller can benefit from a deviation is by cutting his opening offer $a_j(0)$.

Other kinds of equilibrium. In the reciprocal pricing mechanism, the probability with which the buyer is pivotal is simple, it is just the probability that exactly one other buyer re-visits websites. This probability has to be high enough to pay for the cost of search. However, once a seller deviates by changing his initial offer, this could, in general, trigger a sequence of asks that depend on the random number of bidders who choose to revisit websites. This section collects a series of results related to these equilibria.

Suppose the probability that a buyer's value is high enough to revisit when he sees a deviation is π. Define

$$b(t, \pi) = \frac{n!}{t!(n-t)!}(\gamma\pi)^t(1 - \gamma\pi)^{n-t}$$

and

$$q(t, \pi) = \frac{(n-1)!}{t!(n-1-t)!}(\gamma\pi)^t(1 - \gamma\pi)^{n-1-t}.$$

These probabilities describe a distribution of random outcomes in which t represents the number of buyers who re-visit websites. The probability $b(t, \pi)$ represents the probability there are t website visits from the perspective of a seller who considers deviation. The probability $q(t, \pi)$ represents the probability that t of the other buyers choose to visit websites from the viewpoint of a buyer who has just observed a deviation.

Observe from this that

$$b(t, \pi) = \frac{n}{n-t}(1 - \gamma\pi)q(t, \pi).$$

Fix a symmetric equilibrium pricing function $a(\cdot)$, and some deviation $a'(\cdot)$. The value $a(t)$ represents the ask the deviator expects the other sellers to make after t buyers re-visit his their websites. The maximum payoff associated with a deviation is relativly easy to find since buyers cannot observe the values of $a'(t)$ when $t > 0$. Since all sellers set the same price, if $a'(t) > a(t)$ then the deviator's ask will always be one of the m highest bids or asks, and the seller will fail to trade. So the seller will never deviate to $a'(t) > a(t)$ if $a(t) > c$. If $a'(t) < a(t)$, then the deviator will sell for sure at price $a(t)$ unless either there are no buyers whose value exceeds $a'(t)$, or at the price $a'(t)$ if there is exactly one buyer whose value is above $a'(t)$. In this latter case the seller will sell at price $a'(t)$. So whatever $a(t)$ is seller will want to set a price below but as close as possible to $a(t)$ if $a(t) > c$, but will set a price above $a(t)$ when $a(t) < c$.

Putting this together, the deviating seller's payoff is then given by (or is arbitrarily close to)

$$\sum_{t=0}^{n} b(t, \pi) \max[a(t) - c, 0]$$

As deviations are unprofitable they must trigger re-visits by buyers. Thus there is a π such that

$$(1.5) \qquad \sum_{t=0}^{n} b(t, \pi) \max[a(t) - c, 0] \le (a(0) - c)(1 - F(a(0))) \frac{n}{m}.$$

Buyers, on the other hand will re-visit websites after a deviation if they believe it will lower the price they will have to pay enough to compensate them for the search cost. This condition is given by

$$(1.6) \qquad \sum_{t=0}^{n-1} q(t, \pi)(\max[v - a(t+1), 0] - \max[v - a(t), 0]) \ge \epsilon.$$

Our first result concerns the price functions $a(t)$ that are most likely to induce buyers to re-visit websites. The implication of the result is that if a rule $a(t)$ deters deviations and some buyer of type v wants to search after seeing a deviation, then the same buyer will want to search is the pricing rule is given by

$$\tilde{a}(t) = \begin{cases} \max_t a(t) & t < t^M \\ 0 & \text{otherwise,} \end{cases}$$

where t^M is the modal number of buyers who search after seeing a deviation.

Lemma 3. *Suppose there is a symmetric equilibrium with $a(0) > c$. Suppose that a deviation that induces $a'(0)$ causes buyers to re-visit websites with probability π. Let t^m be the modal value of the binomial distribution associated with success probability π. Let $q(t, \pi)$ be the probability of t successes associated with this distribution. Then for any buyer v for which (1.6) holds,*

$$q\left(t^M, \pi\right)\left(v - \max\left[v - \max_t a(t), 0\right]\right) \geq \epsilon.$$

Proof. The gain (or loss) to revisiting websites when t other buyers choose to visit is $\max\left[v - a(t+1), 0\right] - \max\left[v - a(t), 0\right]$. Suppose to start that $v \leq \max_t a(t)$. Consider the problem

$$\max_{a_0,\ldots a_n} \sum_{t=0}^{n-1} q(t, \pi)(a_t - a_{t+1})$$

subject to the constraint that $0 \leq a_t \leq v$ for each t. Note that under these constraints, $v \geq a_t$. The by (1.6) the solution to this problem must yield value at least ϵ.

If we let λ_t and β_t be multipliers associated with the constraints $a_t \leq v$ and $a_t \geq 0$ respectively, then the first order condition for a_t is given by

$$b(t, \pi) - b(t-1, \pi) - \beta_t + \lambda_t = 0.$$

The terms $-\beta_t$ and λ_t have opposite signs. By the properties of the binomial distribution

$$b(t, \pi) - b(t-1, \pi) > 0$$

for t less than or equal to the modal value $\pi\gamma(n-1)$, while it reverses sign and is negative when $t > \pi\gamma(n-1)$. The first order condition then requires that $\lambda_t \neq 0$ for $t < \pi\gamma(n-1)$ and $\beta_t \neq 0$ if $t > \pi\gamma(n-1)$.where the β_t and λ_t are non-positive. Complementary slackness then requires that $a(t) = v$ for each t less than or equal to the modal value, with $a(t) = 0$ otherwise.

If $v > a(t)\,\forall t$, then repeat the argument with the constraint that $a_t \leq \bar{a}$ for each t. $\square$

What Lemma 3 shows is that if there is a symmetric equilibrium with $a(0) > c$, then any buyer type v who was willing to re-visit websites after a deviation would also be willing to revisit websites if the pricing function were given by

$$(1.7) \qquad\qquad a'(t) = \begin{cases} \max_t a(t) & t \leq t^m \\ 0 & \text{otherwise} \end{cases}$$

provided buyers continued to visit websites with probability π.

Heuristically, the best way to induce buyers to revisit is to offer them a pricing rule that doesn't respond at all to buyer re-visits until exactly $t^m + 1$ buyers visit, at which point ask prices are cut to 0. For brevity, we refer to a pricing rule like $a'(\cdot)$ as a *trigger rule*. A trigger rule is characterized by an ask price $a(0)$ and the number of website re-visits that buyers have to make in order to trigger the price cut. For example, the rule given by (1.1) in the first part of this paper is a trigger rule with the price cut triggered by 2 revisits.

The trigger rule $a'(\cdot)$ given by (1.7) won't necessarily be an equilibrium just because $a(0)$ is supported as an equilibrium. Deviating firms will often receive higher prices which may more than make up for the lower prices they receive after a price cut.

Also, notice that if a buyer of type v is willing to re-visit websites when other buyers re-visit with probability π and the pricing rule is $a'(\cdot)$, then every higher buyer type is also willing to re-visit. The value of π is the measure of the set of buyer types who choose to re-revisit in the original equilibrium, and the set who wants to visit under the rule a' could be larger. So it isn't immediate that π is part of a continuation equilibrium with respect to a'.

The probability with which buyers believe their website revisits will lead to a price cut under a trigger rule that cuts price after k visits is given by

$$(1.8) \qquad \frac{(n-1)!}{(k)!\,(n-1-k)!}\,(\gamma\pi)^k\,(1-\gamma\pi)^{n-1-k}.$$

Heuristically, the way to support a trigger rule as an equilibrium is to make the trigger probability (1.8) as high as possible. The best way to accomplish this is to set π so that the number of visits k in the formula above is the modal value of the binomial distribution of visits. As described above, this is accomplished by setting π equal to $\frac{k}{\gamma(n-1)}$.

Lemma 4. *For any k*

$$\frac{(n-1)!}{k!\,(n-1-k)!}\left(\frac{k}{n-1}\right)^k\left(1-\frac{k}{n-1}\right)^{n-1-k} \geq$$

$$\frac{(n-1)!}{(k+1)!\,(n-k-2)!}\left(\frac{k+1}{n-1}\right)^{k+1}\left(1-\frac{k+1}{n-1}\right)^{n-k-2}$$

Proof. Note that

$$\frac{(n-1)!}{k!\,(n-1-k)!}\left(\frac{k}{n-1}\right)^{k}\left(1-\frac{k}{n-1}\right)^{n-1-k}=$$

$$\frac{(n-1)!}{k!\,(n-1-k)!}\frac{n-k-1}{k+1}\left(\frac{k+1}{n-1}\right)^{k+1}\left(1-\frac{k}{n-1}\right)^{n-2-k}=$$

$$\frac{(n-1)!}{(k+1)!\,(n-1-2)!}\left(\frac{k+1}{n-1}\right)^{k+1}\left(1-\frac{k}{n-1}\right)^{n-2-k}\geq$$

$$\frac{(n-1)!}{(k+1)!\,(n-1-2)!}\left(\frac{k+1}{n-1}\right)^{k+1}\left(1-\frac{k+1}{n-1}\right)^{n-2-k}$$

$\square$

The implication of this lemma is that if some type v is willing to re-visit websites when the seller is using a trigger rule that cuts prices on the k^{th} website visit, then the same buyer will be willing to re-visit websites if the seller uses a trigger rule that cuts prices on the second website visit. However, in this case, this is only true if the probability with which buyers re-visit websites is adjusted so that the modal value of the binomial distribution is 1 instead of k.

Equilibrium in Large Markets.

Now we come to main question in the paper. Sellers need buyers to communicate deviations to support equilibrium with non-competitive prices. If it is very costly for buyers to communicate, they won't, and the double auction will provide efficient outcomes. In modern digital markets, communication is relatively cheap, but not free. We are trying to get a handle on how large these costs have to be before sellers lose their ability to exploit this communication.

As in the first section of the paper, we can write down sufficient conditions for equilibria to be sustained with price above marginal cost. Yet it is hard to see from these conditions what magnitudes the relevant costs have. We'll assume markets are large here so we can use limit results to make these more intuitive.

To keep from being bogged down dealing with arcane off equilibrium beliefs, we'll assume that when there is a deviation by some seller, buyers will only change their beliefs about the pricing rule being offered by the deviating seller. Their beliefs about the pricing rules be used by non-deviators should remain unchanged. This restriction is fairly

standard[10] but is not required by weak perfect Bayesian equilibrium. Finally, in all the theorems that follow, we restrict attention to symmetric equilibrium outcomes in which all firms use the same mechanism and all buyers use the same reporting and bidding rules.

The main results of the paper are given in the following theorem.

Theorem 5. *For very large n, a necessary condition for the existence of an equilibrium in which the on path trading price is $p > c$ is*

$$(1.9) \qquad \lim_{n \to \infty} \frac{\epsilon(n)}{v^h} \le e^{-1}$$

where v^h is the maximum value in the support of F. A sufficient condition for the existence of such an equilibrium is that (1.9) holds with strict equality and

$$(1.10) \qquad 2 \lim_{n \to \infty} \frac{\epsilon(n)}{p} \le (1 - F(p)) \lim_{m,n \to \infty} \frac{m}{n}.$$

The upshot of the theorem is that inefficient outcomes can only arise if communication costs for buyers are no more than about $\frac{1}{3}$ of the highest value any buyer assigns to the good. The somewhat surprising element to this is that ease of communication is supposed to be one of the benefits of modern digital markets. Cheap communication is what makes it possible for consumers to discover all the alternatives in a market. However, this same ease of communication also makes it possible for consumers to convey additional information as they do their market discovery.

The first part of the proof is concerned with showing that the efficient equilibrium in which all buyers and sellers bid non-cooperatively in the double auction is unique when costs are large enough (in particular, when they are larger than e^{-1} times the price). In outline, the proof exploits the logic in Lemmas 3 and 4 to show that if any equilibrium in which $p > c$ can be supported, it can be supported with a trigger rule like the one described by (1.1) . It then shows that if the condition in the theorem fails, a trigger rule like (1.1) cannot be supported as an equilibrium because buyers cannot be counted on to convey their market information.

The second part of the theorem puts a bound on how high prices can go. If the price sellers are trying to support is too high, no buyers will be willing to pay it. The basic argument in the proof is to show that the probability that a deviator an cut price and get away with it is $2 \lim_{n \to \infty} \epsilon(n)$ when sellers use pricing rules like (1.1). In that case the

[10]For example, it is given as part of the definition of perfect Bayesian equilibrium by (Osborne and Rubinstein 1994).

profit of a seller is $(p - c)$. If a seller decides not to deviate, his profit on the equilibrium path is the probability he sells at p times $(p - c)$. The selling probability is $(1 - F(p)) \lim_{n \to \infty} \frac{m}{n}$.

Proof of the Necessity of (1.9) in Theorem 5.

Proof. Let $a(t)$ be a pricing rule that is supported in equilibrium, and let π be the probability with which buyers re-visit websites after a deviation. In this argument, it doesn't matter what the deviation is, or whether the re-visiting probability depends on this deviation. Let v be the lowest value buyer who revisits websites in the continuation equilibrium. By the continuity of (1.6) (the expected payoff associated with re-revisiting websites) in v, (1.6) holds with equality for a buyer of type v.

Then by Lemma 3, the trigger rule (1.7) satisfies

$$\frac{(n-1)!}{(k)!(n-1-k)!} (\gamma\pi)^k (1 - \gamma\pi)^{n-1-k} \geq \frac{\epsilon(n)}{\min[a'(0), v]},$$

while by Lemma 4,

$$(n-1)\gamma\pi (1 - \gamma\pi)^{n-2} \geq \frac{\epsilon(n)}{\min[a'(0), v]}.$$

The left hand side of this last expression attains it maximum when $\pi = \frac{1}{\gamma(n-1)}$, at which point it is equal to $\left(1 - \frac{1}{n-1}\right)^{n-2}$ while the right hand side is always at least as large as

$$\frac{\epsilon(n)}{v^h}$$

so price cuts cannot be triggered if

$$\left(1 - \frac{1}{n-1}\right)^{n-2} < \frac{\epsilon(n)}{v^h}.$$

Taking limits gives the necessary condition in the theorem. So if the condition in the theorem fails, there cannot be an equilibrium with $a(0) > c$ when n is large. $\qquad\square$

A few comments are in order. First, note that the probability $\gamma(n)$ doesn't appear anywhere in the set of sufficient conditions. The reason for this is that we assumed from the outset that $\pi\gamma(n) \geq \frac{1}{n-1}$. It is possible to extend the conditions above to handle environments where this condition fails. When it does, the upper bound on buyer communication costs (above which collusive equilibrium can no longer be supported) varies with γ in a complex way. Since this problem is

only interesting when there is some significant probability that buyers observe some seller commitments, we leave this part out.

Proof of the Sufficiency of (1.10) in Theorem 5.

Proof. Suppose all firms use rules like (1.1). By (1.10), $p < v^h$ (recall v^h is the largest element in the support of F). Therefore, by (1.9), there is an interval $[v^*, v^h]$ for which $e^{-1} > \frac{\epsilon(n)}{p} > \frac{\epsilon(n)}{v}$ for every $v \in [v^*, v^h]$. The function

$$(n-1)\gamma(1 - F(v))(1 - \gamma(1 - F(v)))^{n-2}$$

is continuous in v on the interval $[v^*, v^h]$. For n very large, the peak of this function is close to e^{-1}, while

$$(n-1)\gamma(1 - F(v^*))(1 - \gamma(1 - F(v^*)))^{n-2}$$

tends to zero with n. It follows by the intermediate value theorem that for large enough n, there is a $\hat{v} < v^*$ and $(1 - F(\hat{v}))\gamma > \frac{1}{n-1}$ at which

$$(n-1)\gamma(1 - F(\hat{v}))(1 - \gamma(1 - F(\hat{v})))^{n-2} = \frac{\epsilon(n)}{p}.$$

Substituting this into the payoff of the deviating seller gives his payoff as $(p - c)$ times

$$(1.11) \qquad (1 - \gamma(n)(1 - F(\hat{v}_n)))^n + \frac{n(1 - \gamma(n)(1 - F(\hat{v}_n)))}{(n-1)}\frac{\epsilon(n)}{p}$$

which is the probability with which his deviation goes unreported.

Since $\hat{v}_n$ must go to v^h, the second term in (1.11) has limit $\lim_{n\to\infty}\frac{\epsilon(n)}{p}$. We want to focus on the first term. This term is

$$(1 - \gamma(n)(1 - F(\hat{v}_n)))^n = (1 - \gamma(n)(1 - \hat{v}_n))^2(1 - \gamma(n)\pi r)^{n-2} =$$

$$\frac{(n-1)\gamma(n)(1 - F(\hat{v}_n))}{(n-1)\gamma(n)(1 - F(\hat{v}_n))}(1 - \gamma(n)(1 - F(\hat{v}_n)))^2(1 - \gamma(n)(1 - F(\hat{v}_n)))^{n-2} =$$

$$\frac{(1 - \gamma(n)(1 - F(\hat{v}_n)))^2}{(n-1)\gamma(n)(1 - F(\hat{v}_n))}\frac{\epsilon(n)}{p}.$$

Now

$$\frac{(1 - \gamma(n)(1 - F(\hat{v}_n)))^2}{(n-1)\gamma(n)(1 - F(\hat{v}_n))}\frac{\epsilon(n)}{p} \leq$$

$$\left(1 - \frac{1}{n-1}\right)^2\frac{\epsilon(n)}{p}$$

because of the fact that $\gamma(n)(1 - F(\hat{v}_n)) > \frac{1}{n-1}$. So the limit of the expression in (1.11) is less than or equal to $\lim_{n\to\infty} 2\frac{\epsilon(n)}{p}$ as given in the theorem.

This gives

$$2 \lim_{n \to \infty} \frac{\epsilon(n)}{p}(p-c) \leq (1 - F(p))(p-c) \lim_{m,n \to \infty} \frac{m}{n}$$

so a deviation is unprofitable.   □

This theorem doesn't give a complete characterization. It doesn't rule out the possibility that prices that can't be supported by reciprocal pricing rules can be supported with other trigger rules, or with rules that don't involve trigger like reactions. It seems unlikely that these alternatives could accomplish this. For example, a trigger rule that cuts prices after $k > 2$ website re-visits requires that buyers choose to re-visit with higher probability. When n is large, whatever this probability is, it supports a distribution of buyer visits for which the probability that $k-1$ buyers visit is close to $\frac{\epsilon(n)}{p}$. A reciprocal pricing rule with the same starting price, supports a distribution that for which the probability that 1 buyer visits is close to $\frac{\epsilon(n)}{p}$, but with a lower variance. As a consequence, the probability that 1 or few buyers re-visit should be smaller than the probability that $k-1$ buyers visit, meaning deviations are less profitable under the reciprocal pricing rule. So far, no proof of this is available.

There are many ways that the assumptions in the model could be relaxed. For example, instead of assuming sellers have the same costs, their costs could be drawn randomly from the distribution F, exactly as buyers values are.

The main complication this generates is that trading prices on the equilibrium path are then random. Whe firms are maintaining a price above the competitive price, there will be a strictly positive probabilility that there will be many high value buyers who would bid the price up above the collusive price in the double auction.

The computation of the probabiity that the buyer is pivotal in instituting the punishment is similar, but for high value buyers the gain to being pivotal and choosing to re-visit becomes a random variable. Though, similar (more complicated) calculations will apply, they will not admit the simple $\frac{1}{e}$ characterization the upper bound on costs in the limit.

1.3. **Conclusion.** The paper shows conditions under which inefficient equilibrium can and can't be supported as equilibrium. This is nicely summarize by looking at the case where markets are very large. Then inefficient high priced equilibrium disappear once the costs of communicating market information to seller reach about $\frac{1}{3}$ of the highest value possessed by any buyer.

There are two parts of the story that differ from traditional mechanism design. The first is that some of the people who participate in a mechanism may not know exactly how the mechanism works. Traders who don't know what a mechanism designer is doing can't give other mechanisms designers any useful information. The second is that buyers can't send complicated messages, and may find it costly to send any messages at all.

As appears to be true in most digital markets, even buyers who can see aspects of the sellers' mechanisms really understand very little about them (though they do guess them correctly in equilibrium).

The logic of the argument resembles the logic in repeated games, though it is simpler in that it doesn't involve any discounting. As with repeated games, the fact that inefficient equilibrium can exist doesn't mean that they do. In many ways the contribution of this kind of argument is to identify simple ways that the equilibria can be supported. Reciprocal pricing contracts are an example of the way that sellers can provide buyers with incentives to enforce their own collusion without requiring any methods that aren't already in use on the internet. A useful direction to push this idea is to look for alternative models of the way that buyers and sellers communicate in the hope that something will suggest a way to search for hard evidence that these equilibria can occur.

In the meantime, the main lesson would seem to be that 'first welfare theorem' like arguments that large markets with lots of competition should promote efficiency should at least be held in abeyance.

Full knowledge of sellers' commitments and costless communication of type are absolutely fundamental assumptions in standard mechanism design. One of the implications of the results here is that it is interesting to examine standard mechanism design to study the case where some agents don't understand the rules that are being used to determine allocations.

As an example, the logic of second price auctions rapidly falls apart if none of the bidders actually knows for sure that the second price is being used. The buyers will sensibly believe that the seller will use first price techniques. In the extreme case in which no buyer understand the rules, a fixed price mechanism is the only feasible mechanism. In the intermediate case in which some buyers do understand the pricing rules, the analysis would probably resemble the analysis in this paper.

References

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

FORGES, F. (2012): "A folk theorem for Bayesian games with commitment," *Games and Economic Behavior*, 78, 64–71.

MARTIMORT, D., AND L. STOLE (1998): "Communications Spaces, Equilibria Sets and the Revelation Principle under Common Agency," University of Chicago unublished manuscript.

MASATLOGLU, Y. (2015): "A Random Inattention Model," Discussion paper, University of Michigan.

MASATLOGLU, Y., D. NAKAJIMA, AND E. OZBAY (2012): "Revealed Attention," *American Economic Review*, 102(5).

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

OSBORNE, M. J., AND A. RUBINSTEIN (1994): *A Course in Game Theory*. MIT PRess, Cambridge Massachusetts.

PECK, J. (1994): "Competition in Transactions Mechanisms: The Emergence of Competition," Ohio State University mimeo.

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

——— (2001b): "Surplus Extraction and Competition," *Review of Economic Studies*, 68(3), 613–633, University of Toronto, manuscript.

——— (2015): "Reciprocal contracting," *Journal of Economic Theory*, 158, Part A(0), 102 – 126.

PETERS, M., AND S. SEVERINOV (1997): "Competition Among Sellers who offer Auctions Instead of Prices," *Journal of Economic Theory*, 75(1), 141–179.

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.

SALOP, S. (1986): *Practises that Credibly Facilitate Ologopoly Coordination*Cambridge, MIT Press.

VIRAG, G. (2010): "Competing auctions: Finite markets and convergence," *Theoretical Economics*, 5, 241–274.

WATSON, J. (2012): "Contractual Chains," Discussion paper, University of California at San Diego.

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.