Dynamic Mechanism Design with Costly Participation

Katalin Bognar, Tilman Börgers, Moritz Meyer-ter-Vehn

February 2008
Preliminary and Incomplete

Abstract

We introduce costly participation into a dynamic mechanism design model. Depending on the assumptions on the timing of participation timing the revelation principle may fail to hold. Dynamic mechanisms generally allow for sequential coordination of participation decisions and may therefore improve upon simultaneous revelation mechanisms. This failure increases the difficulty of finding optimal mechanisms.

We apply this idea to an auction setting and discuss how the optimal auction depends on the specification of the participation timing. In particular we find that an open out-cry auction that allows potential bidders to observe others' bids is revenue-optimal in a large class of auctions.

1 Introduction

In many applications of mechanism design theory agents bear costs from participating in the mechanism. Bidders may need to pre-qualify for an auction, supply certification of funds and travel to the auction site. The same may be true for prospective buyers and sellers in a marketplace exchange. In a committee voting setting, committee members may be required to attend long meetings before being allowed to cast their vote.

Most of the mechanism design literature ignores these costs, assuming that agents care only about the result of a mechanism, e.g. allocation of goods and monetary transfers, or the result of a vote. In contrast, they are assumed to attach no value or cost to any action taken in the mechanism or to the mechanism per se. In this paper we propose a dynamic mechanism design model that captures participation costs and study optimal mechanisms in this model.

It is clear that costs of participation provide a rationale for reducing participation in a mechanism to a minimal level required to reach the desirable outcomes. When a strong bidder has entered an auction it is of little benefit to anybody if a weaker bidder incurs costs to enter as well. Analoguely if the outcome of a dynamic voting procedure seems sufficiently clear after half of the committee members have cast their votes it may not be worthwhile for the remaining members to show up. From these examples it becomes clear that there may be benefits from deferring participation decisions of some agents (e.g. auction entry of weak bidders) from the beginning of the game to a later stage when they can be made contingent on actions of other agents.

This option value of deferring the costly entry decision will be a major determinant for optimal mechanisms. Thus, a key modeling assumption in this paper will be how far participation can be deferred or, put another way, what actions are allowed to happen and what an agent is allowed to learn before she takes her participation decision. Given the breadth of mechanism design applications and the different interpretations of the participation costs as travel and attendance costs, certification costs, funding costs, etc., it would be unreasonable to expect one assumption to be the most reasonable for all cases. Therefore we will discuss four alternative assumptions on what can happen before participation.

"Simultaneous participation", the most restrictive one, requires all agents to take their participation decisions simultaneously at the beginning of the game before they had the chance to learn anything about others. "Sequential participation" allows for sequential participation where an agent can learn nothing about past play before deciding to participate in the mechanism. "Activating participation" allows for more sequentiality in terms of "wait-and-see" strategies that allow an agent to pass a participation node and get another chance to participate later in the game. "Deferred participation", the least restrictive one, puts no restrictions at all on what an agent can do in a game before she participates.[1]

While we consider "activating participation" to be the most reasonable of the above assumptions, this general approach of considering also the other set of assumptions allows us to nest models used in the literature in a larger framework and highlight how the differences in the assumptions lead to different features of the optimal mechanisms.

It is common in the mechanism design literature to invoke the revelation principle to reduce the analysis to direct mechanism: Instead of having agents perform their actions according to some equilibrium strategies in some indirect mechanism, the designer asks every agent for her information and then instructs her what actions to take in the ensuing game. The second stage, of actually playing the game, is then usually cut short by having the designer impose the equilibrium outcomes directly.

The problem with this approach in our setting is that there is no satisfactory interpretation of participation costs (an applied feature) in a direct mechanism (a theoretical concept). Formally, the question is how the revelation stage relates in a chronological order to the participation decision. Taking a narrow stance on this by requiring the participation decision to happen before the revelation stage is tantamount to interpreting the participation costs as communication costs. This seems unreasonable when an agent's private information is summarized, say, by a single number describing her valuation of an object to be auctioned.[2] Allowing for type revelation before the participation decision, on the other hand, leads to the paradoxy that the instructed "participation decisions" are actually the last actions in the mechanism.

It should therefore come at no surprise that for "activating participation", arguably the most interesting of the four specifications introduced above, there is no revelation principle, in the sense that there is no natural set of direct mechanisms that implements the same choice rules as the indirect mechanisms.[3] We would like to emphasize that we see this as a virtue of this specification[4] rather than a shortcoming, as it allows a role for dynamic mechanisms that allow agents to coor-

[1] Nevertheless, feasible outcomes may be restricted by participation. E.g. in a procurement auction the chosen supplier may need to undergo a costly certification process. However, the auctioneer may in many cases defer this costly "participation" until after the auction to make only the winner incur these costs.

[2] Note that communication costs may well be relevant when considering large mechanisms, like the determination of equilibrium prices for the overall economy.

[3] On the other hand we do find revelation principles that relate indirect mechanisms with "simultaneous participation" to direct mechanisms where type revelation is costly, i.e. happens after the participation decision, and indirect mechanisms with "deferred participation" to direct mechanisms where type revelation is not costly, i.e. happens before the participation decision.

[4] and more generally of all specifications that are less restrictive than "simultaneous participation" but more restrictive than "deferred participation"

dinate their entry decisions endogenously. This relates to dynamic auctions observed in practice that allow weak bidders to see whether a strong bidder has entered, before making the costly entry decisions herself. Alternatively, it relates to dynamic voting procedures that allow voters with weak preferences/information to first see whether the vote is close, before incurring the voting costs.

Another intriguing consequence of this failure of the revelation principle is that it becomes unclear how to generally identify optimal mechanisms. The standard approach for doing so is to maximize the objective function over directly implemented choice rules, given the constraints of incentive compatibility, individual rationality, and sometimes, budget balance. Here, we face an extra constraint "participation feasibility", namely existence of an indirect mechanism that implements the desired choice rule and the desired participation rule simultaneously. Surprisingly we are still able to find optimal and efficient auction. The reason we are able to do so is that the "participation feasibility" constraint is not binding, i.e. is satisfied by an optimal unconstrained mechanism.

Applying the model to an auction setting, our main result is that a standard open out-cry auction with an appropriate reserve price (resp. no reserve price) is the revenue-optimal (resp. efficient) mechanism. This is remarkable because this auction satisfies "activating participation" - in that a player may observe the auction for free but needs to incur the participation costs to place a bid - yet it is easily seen to be optimal in the class of mechanisms with "deferred participation" - i.e. auctions where bidders can bid for free and only the winner needs to incur the participation costs ex-post. We note that the open out-cry auction is strictly superior to any mechanism with "simultaneous participation".

Let us briefly relate the participation costs discussed here to other costs in the mechanism design literature and highlight the differences. "Information acquisition costs", such as in Bergemann and Valimaki (2002) model the direct costs and opportunity costs of learning one's own valuation and other information. They differ fundamentally from the costs considered here in that they are incurred before the agent knows her valuation.[5]

"Search costs", such as in McAfee, McMillan (1988) are costs borne by the designer, e.g. the government when identifying potential suppliers for a procurement auction, that are also incurred in ignorance of the agent's valuation. As agents are typically ex-ante symmetric there is no role for endogenous coordination of entry in these models.

"Cognitive costs" are costs borne by the agent to learn about the mechanism and past play. Our model does not cover these as it imposes costs only on certain actions rather than on, say, the number of nodes reached or the complexity of the game tree.

A common consequence of all these cost types is that they introduce an option value of delaying participation and thus a rationale for dynamic mechanisms.

Before going into the literature review, we would like to point out one major restriction of our model: Participation costs are assumed to be equal across mechanisms. Thus, in an auction setting, the potential buyer of some object incurs the same costs from bidding in a time-consuming auction as she does from accepting a "take-it-or-leave-it" offer. While this is clearly an abstraction from reality, this assumption is necessary to ensure that a mechanism cannot be optimal because it has exogenously lower costs or agents simply like it.

[5]Note however that the differentiation between these two categories may be less clear than it seems at first sight: Lengthy committee meetings to decide about a course of action to reach a common goal could be viewed either simply as opportunity costs or as the sharing, or the mutual acquisition, of each others information.

1.1 Literature Review

Participation Costs

While there is a large literature on mechanism design with costs related to participation, none of it is truly similar in spirit to this paper. Samuelson (1985) notes that restricting entry may increase both efficiency and revenue in a 1^{st} price auction with simultaneous costly participation, essentially by reducing the chance of miscoordinated entry. More generally, Stegemann (1996) and Tan and Yilankaya (2006) identify the efficient mechanism while Celik and Yilankaya (2007) identify the revenue-optimal auction. Both of these papers focus on mechanisms with simultaneous participation and do not consider the possibility of dynamic coordination of participation.[6] Closer to our model is Ehrmann and Peters (1994) who find that "fix price of best offer" is an optimal selling mechanisms when buyers are contacted sequentially in a random order and cannot learn anything about past play before taking the costly participation decision.[7] This corresponds to our specification of "sequential participation". Thus, our framework nests the above models, and highlights how assumptions on the timing of participation shape the features of optimal mechanisms.

Börgers (2004) studies a voting model where voting is costly and finds that under certain circumstances sequential voting dominates simultaneous voting. We plan to apply our framework to a voting setting and to go beyond this result in allowing for endogenous order and asking for the optimal voting rule, rather than comparing two individual institutions.

Search Costs

McAfee and McMillan (1988) and Cremer, Spiegel and Zheng (2006) identify optimal search mechanisms (where searching is costly for the designer but participation is free for the agents). These settings admit revelation principles reducing the set of all imagineable search mechanisms to direct mechanisms with the property that each agent is contacted at most once and in an order that is either arbitrary (in McAfee, McMillan) or determined by the ex-ante asymmetry of the bidders (in Cremer, Spiegel, Zheng). The principle pertains essentially because mechanisms are assumed to be designer-centric leaving no room for endogenous coordination between the agents. Whereas the optimal mechanism in McAfee and McMillan (1988) is "fix price or best offer", Cremer, Spiegel and Zheng (2006) find additional effects stemming from the asymmetry of bidders.

Information Acquisition Costs

Bergemann and Valimaki (2002) consider a general quasi-linear mechanism design setting in which agents simultaneously invest into learning the state of the world before participating in a revelation mechanism. They show that investment will be (constrained) efficient when valuations are private. Zheng (2005) generalizes this result to dynamic mechanisms where the investment decisions can be contingent on the results of own and others' previous learnings. These results follow from the basic insight that agents capture their informational rent in Groves mechanisms and therefore have the correct incentives to invest in information.

McAfee and McMillan (1987) study auctions with costly information acquisition and find that the first price auction without reserve price is a (constrained) optimal mechanism. Although their model is not explicitly dynamic, the argument assumes that bidders enter the auction sequentially

[6] This is surprising in so far as that a simple mechanism like sequential take-it-or-leave-it offers to the potential buyers is easily seen to outperform the constrained optimal mechanism identified by Stegemann:

For two buyers with valuations uniform i.i.d. on $[1;2]$ and participation costs of 0.3, Stegemann shows that the asymmetric equilibrium of the 2nd price auction is constrained efficient with a surplus of 1.22 whereras, sequential take-it-or-leave-it offers at $p_1 = 1.2$ and $p_2 = 0.7$ are easily seen to yield a surplus of 1.33.

[7] It would be nice to see how this result changes when bidders are allowed to observe past play. "Fix price or best offer" would certainly be dominated by some mechanism with a decreasing price sequence but it is not cleat what would be optimal.

until the expected profit of an additional entrant falls short of the information acquisition costs. Given this zero-expected profit condition on the bidders the seller will reap all the surplus of the auction, immediately yielding the result that it is not in her interest to reduce this surplus through a reserve price which would deter entry.[8] [9] In a more general setting Bulow and Klemperer (2007) compare an auction (with sequential participation but no early bids) to a sequential bidding process which differs from the auction by allowing for early, pre-emptive bids. As one would expect, the sequential bidding process yields higher social surplus than the auction as it is better at coordinating entry decisions. However, Bulow and Klemperer show that early bidders are able to appropriate enough of this surplus to make sequential bidding revenue-inferior to the auction.

Cremer, Spiegel and Zheng (2007) go beyond both of the above papers by solving for the optimal auction with sequential costly information acquisition.[10] The result is driven by the insights that the optimal auction needs to be efficient and the efficient auction is implementable as agents reap their informational rents in Groves mechanisms.

Gershkov and Szentes (2007) consider a voting setting with common values where jurors have an incentive to free-ride on each other's investment in information and show how an optimal voting scheme addresses this issue by concealing past play from jurors.

To summarize, the literature on search costs and information acquisition costs has gone a long way to study sequential mechanism design[11]. This paper tries to do the same for participation costs. More specifically it contributes to the literature in two ways: First it introduces a dynamic mechanism design model that allows agents not only to enter sequentially but also determine the order of entry in an endogenous way. Second it applies this dynamic model to a standard independent private values auction model. Despite the lack of a revelation principle, we identify the socially efficient and the optimal auction in these settings. Whereas the socially efficient mechanism implements the first-best choice rule, the optimal auction is identified by a generalization of the revenue equivalence principle.

The paper is structured as follows: Section 2 introduces the formal model of mechanism design with participation costs. Section 3 applies the model to a private value auction setting and identifies the optimal and the efficient auctions. Section 4 concludes.

---

[8] The "no reserve price result" is a special case of the "efficiency principle" (see e.g. Cremer, Spiegel and Zheng (2007)) stating that if agents have no private information a priori and the designer controls their access to the information gathering device (but not the realized information) he can extract any agent surplus through entry fees. Therefore the optimal mechanism is an efficient and fully-extracting mechanism.

In contrast to this ex-ante interim individual rationality argument, the participation decision is taken after agents know their valuations (i.e. at the interim stage) when looking at search costs or proper participation costs. Therefore, the revenue-oriented designer faces the classical trade-off between making some profit from the low types and extracting surplus from the high types.

[9] It is worth noting that the first price auction may not be the unconstrained optimal mechanism as the paper implicitly assumes a strict order between entry decisions and the bidding phase. It should therefore come at no surprise that a take-it-or-leave-it price may do better than the auction because it gives potential bidders the opportunity to make their entry decision contingent on the type realization of a previous entrant, who may have already bought the object making the entry decision obsolete.

Elementary calcuations show that this will be the case for valuations i.i.d. on $[0;1]$ and a) two bidders and costs $c = \frac{1}{6}$ or b) a sufficiently large number N of potential bidders and any costs.

[10] The sequential nature of the optimum is supported by Compte and Jehiel (2008) who show that a dynamic ascending auction will outperform a static 2nd price auction when only one bidder can learn something about his valuation (resp. some bidders can learn their valuation). This is essentially because the bidders can make their information investment contingent on ohters' information.

[11] While the papers cited above are only concerned with optimal search auctions but not the efficient search auction, the latter should be easy to identify with the same methods.

2 Games and Mechamisms with dynamic costly participation

2.1 Participation in choice rules

Consider a standard social choice model where agents $i \in \mathcal{N}$ hold private information $(\theta_i)_{i\in\mathcal{N}} \in \Theta$ with probability μ and a social planner wants to base a **decision** $k \in \mathcal{K}$ on this information according to a choice rule $\psi : \Theta \to \mathcal{K}$. Going beyond the standard model assume that it is costly for an agent to participate in the choice, and to reveal her information. Therefore, we augment the choice rule by a participation rule that keeps track of participants in the mechanism, i.e. we consider choice rules $(\psi, p) : \Theta \to \mathcal{K} \times \{0; 1\}^N$ with the interpretation that agent i participated iff $p^i = 1$. We call $(\psi(\theta), p(\theta))$ the **outcome** of the choice.

In the auction setting we will restrict allocation of the object to participants of the auction. We will thus generally assume that for each subset of participating agents $\mathcal{I} \subseteq \mathcal{N}$ there is a subset of feasible decisions $\mathcal{K}(\mathcal{I}) \subseteq \mathcal{K}$ and focus on (participation) **feasible choice rules** (ψ, p) with the property that $\psi(\theta) \in \mathcal{K}(\mathcal{I})$ requires $p^i(\theta) = 1$ for all $i \in \mathcal{I}$.

2.2 Participation in dynamic games

We now want to apply the classical mechanism design approach to choice settings where participation is costly. We will see below that in this case static mechanism are restrictive and therefore now formalize what it means for some player to participate in a dynamic game.

Definition 1 *For an extensive form game G a* ***dynamic participation structure*** *P consists of a subset of binary action nodes, the "participation nodes" at which the respective player can decide to "participate" or to "not (yet) participate" in the game, with the property that each player can take at most one "participation" decision on any game path.*

Note that each strategy profile π (and thus each terminal node of the game) determines whether agent i participated, $P^i(\pi) = 1$, or not, $P^i(\pi) = 0$.

So far the definition is almost vacuous. We will now introduce the various restrictions on the dynamic participation structure P discussed in the introduction.

Definition 2

1. *Any (G, P) is said to satisfy "deferred participation".*
2. *(G, P) satisfies "Activating Participation" iff: an action node n of player i is a participation node if and only if player i has not yet taken a "participation action" on the game path leading to n.*
3. *(G, P) satisfies "Simultaneous Participation" iff: all players simultaneously decide whether to participate or not before any other action can be taken or the game can end.*
4. *(G, P) satisfies "Sequential Participation" iff: the game is anonymous and each player has a unique participation node (but the game can end before all the participation nodes are reached).*

Note that $3, 4 \Longrightarrow 2 \Rightarrow 1$.

"Deferred participation" may seem like an meaningless definition at first. The only restriction imposed by this specification comes from the fact that non-participation $P^i(\pi) = 0$ of player i

will restrict the set of feasible outcomes $\mathcal{K}(\mathcal{N} \setminus \{i\}) \subseteq \mathcal{K}$. However, consider a large interpretion of "participation costs" as including, say, the costs in a procurement auction of proving compliance with all the product specifications. It is reasonable to allow for these costs to be incurred after the bidding is over and to impose them only on the winner of the auction. While it runs counter to the literal meaning of "participation" to defer it until the end of the game rather than insisting that a player needs to participate before she can take any other action, this specification will serve as a useful benchmark below.

"Activating participation" arguably captures best the nature of "participation". Before a player can take any other action in a game she first has to participate in the game. Note that this leaves open the possibility for a player to observe and react to other players. She can even do so iteratively, by choosing to "wait and see" at some early participation node and then take a participation decision later on the game path, possibly after observing others' actions. This specification stresses the point that acting is costly while observing is not.

"Simultaneous participation" is more restrictive in the sense that it precludes a player from learning anything about the actual play of the game before deciding to participate or not. Thus it models situations in which the participation stage can be completely separated from the actual play of a game, e.g. an auction that requires physical attendance or another form of uncoordinated enrollment of all bidders before the actual bidding starts.

"Sequential participation" models a situation where a mechanism designer contacts the players in a random order and gives them the opportunity to participate. In doing so she is not allowed to convey any information about the past actions of other players. The anonymity assumption ensures that a player cannot infer any information from the fact that she has been contacted (instead of the game having ended before this could happen). While the anonymity assumption may feel unnecessarily restrictive we maintain it to remain consistent with the model in Ehrmann and Peters (1994).

The following lemma describes a sense in which the above definitions (2,3,4) do not allow a player to influence the outcome of the game without participating.

Lemma 3 *Let (G, P) satisfy activating participation. For given strategies π^{-i} of the players other than i, there is exactly one terminal node that can be reached by player i by choice of π^i unless she participates in (G, P).*

Proof. Consider a game path on which actions of others are determined by π^{-i} and on which i does not participate. By definition the only action nodes of i that are reached by the game path are participation nodes at which i's action is assumed to be "not participate". ■

Note that a player can influence the outcome of the game by deciding not to participate in the sense that this non-participation leads to a different outcome than any participation decision. The point of the above lemma is that it is not possible to differentiate between different outcomes given that the player does not participate.

2.3 Mechanisms with costly participation

The last missing ingredient for mechanism design are agents' preferences. So, let agent i's preferences over outcomes (k, p) be described by a utility function $v^i(k, \theta) - p^i c$ where $v^i(k, \theta)$ is i's valuation of outcome k given types θ and p^i indicates whether she participated and c is the costs of participation (equal for all game forms).

Definition 4 *Given information structure (Θ, μ), valuations u and participation costs c a* ***mechanism*** $m = (G, P, \phi)$ *is defined by a game form G with a dynamic participation structure P, and a mapping from terminal nodes to outcomes $\phi : T \to \mathcal{K}$.*[12]

The mapping from terminal nodes to outcomes $\phi : T \to \mathcal{K}$ is restricted to take values in $\mathcal{K}(\mathcal{I}) \subseteq \mathcal{K}$ if only players $i \in \mathcal{I} \subset \mathcal{N}$ participated in the game.

We say that a mechanism $m = (G, P, \phi)$ satisfies activating (resp. deferred, simultaneous, sequential) participation if the game (G, P) does.

As usual we will say that a choice rule (ψ, p) is **implemented** by some mechanism $m = (G, P, \phi)$ if the game induced by the mechanism and the valuation structure (Θ, μ, u, c) has an sequential equilibrium $\pi = \left(\pi^i\left(\theta^i\right)\right)_i$ that leads to the same outcomes as the choice rule, i.e. for all i and $\theta = \left(\theta^i, \theta^{-i}\right)$

$$\begin{aligned} \psi(\theta) &= \phi(\pi(\theta)) \\ p^i(\theta) &= P^i(\pi(\theta)) \end{aligned}$$

We say that a choice rule (ψ, p) is implementable with activating (resp. deferred, simultaneous, sequential) participation if there exists a mechanism $m = (G, P, \phi)$ that implements (ψ, p) and (G, P) satisfies the respective concept.

2.4 Direct Mechanisms and the Revelation Principle

A large part of the mechanism design literature focuses on direct mechanisms in which each agent has a single action node at which she reports her information to the mechanism designer who subsequently implements the outcome specified by the mechanism. By the revelation principle a choice rule that is implementable by any mechanism if and only if it is implementable by a direct mechanism.

While participation costs seem to be relevant in many applications of mechanism design it is far from clear how to interpret them in direct revelation mechanisms where the only action of any agent consists in revealing her information. Formally there are two ways of defining participation in direct mechanisms.

Either one can interpret the announcement of one's type as an action that must be preceded by a participation decision. According to this view, a *direct mechanism with simultaneous participation* consists of an initial participation stage, where agents simultaneously decide whether or not to participate in the mechanism, and an announcement stage where all participants announce their types. The outcome function ϕ then maps participation decision and reports to outcomes $\phi : \prod(\Theta_i \cup \{\varnothing\}) \to \mathcal{K}$. Following Stegemann (1996) we call a choice rule (ψ, p) *(simultaneously) incentive compatible* iff the respective direct mechanism with simultaneous participation has an equilibrium where agent i participates iff $p^i\left(\theta^i\right) = 1$ and every participating type reports her type θ^i honestly $\widehat{\theta}^i\left(\theta^i\right) = \theta^i$.

Alternatively one can interpret the type announcements as free preplay communication. According to this view, a *direct mechanism with deferred participation* consists of an initial communication stage, where all agents simultaneously announce their types, followed by a participation stage in

[12] As the game form G defines a mapping from strategies π to terminal nodes t, we will abuse notation to also allow strategies π as arguments for ϕ.

which the designer instructs the relevant agents to "participate" in the mechanism. Again we call a choice rule (ψ, p) *(deferred) incentive compatible* iff the respective direct mechanism with deferred participation has an equilibrium where agent i reports her type θ^i honestly $\widehat{\theta}^i(\theta^i) = \theta^i$ and every agent obeys the designers instructions whether or not to participate (when $p^i(\theta) = 1$).

We are now ready to state a major conceptual difference between dynamic mechanism design with participation costs and the classical case (such as in Myerson (1981) or Stegemann (1996)).

Proposition 5 (Revelation Principle) *1. **Simultaneous participation:** A choice rule (ψ, p) is implementable by a mechanism with simultaneous participation iff it is simultaneously incentive compatible.*

*2. **Deferred participation:** A choice rule (ψ, p) is implementable by a mechanism with deferred participation iff it is deferred incentive compatible.*

*3. **Failure of revelation principle for activating participation:** There are choice rules (ψ, p) that are implementable by mechanisms with activating (resp. sequential) particpation but are not simultaneously incentive compatible. Conversely, there are deferred incentive compatible choice rules (ψ, p) that are not implementable by mechanisms with activating (resp. sequential) participation.*

Proof. The proof of part 1) and 2) is identical to the proof of the standard revelation principle, and is therefore omitted.

Part 3) is proven by example. Consider a mechanism with activating participation that lets agents participate sequentially and ends as soon as one agent chooses to participate (e.g. sequential take-it-or-leave-it offers to buy some object). The participation decision of agent i in such a mechanism will generally depend on others' actions, and thus types, i.e. $p^i(\theta)$ depends not only on θ^i but also on θ^{-i}. In contrast, we have $p^i(\theta) = p^i(\theta^i)$ for any simultaneously incentive compatible choice rule.[13]

Conversely, observe that by Lemma 3 any choice rule (ψ, p) that is implemented by a mechanism with activating participation must lead to the same outcome $\psi(\theta^i, \theta^{-i}) = \psi(\theta^{i\prime}, \theta^{-i})$ if we hold other agents' types θ^{-i} fixed and neither type of i participates $p^i(\theta^i, \theta^{-i}) = p^i(\theta^{i\prime}, \theta^{-i}) = 0$. Clearly, not every deferred incentive compatible choice rule needs to satisfy this property. To the contrary, in a direct mechanism with deferred participation agent i can report his type, and thus differentiate between θ^i and $\theta^{i\prime}$ without participating. ■

Without the revelation principle it is unclear how to characterize mechanisms that are optimal with respect to some objective function (such as social surplus, revenue or any other objective) as the set of possible mechanisms is hard to describe.

Denote by CR^f the set of feasible choice rules, and by CR^{act} (resp. $CR^{sim}, CR^{def}, CR^{seq}$) the set of choice rules that are implementable by a mechanism with activating (resp. simultaneous, deferred, sequential) participation. By the above we have $CR^{sim}, CR^{seq} \subsetneq CR^{act} \subsetneq CR^{def} \subseteq CR^f$.

2.5 Evaluation of choice rules

Rather than characterizing further the set of all implementable choice rules, we will focus in the next sections on identifying implementable choice rules that are optimal for some preference of the

[13] It seems that the revelation principle does not even hold for mechanism design with a single agent: A general mechanism can start with a random move subjecting the agent to play in either one of two subgames, in which different types may want to participate. This kind of result cannot be replicated by a static revelation mechanism.

mechanism designer.

Let the preferences of the designer given outcome $\left(k,\left(p^i\right)_i\right)$ in state θ be given by $S : K \times \{0;1\}^N \times \Theta \rightarrow \mathbb{R}$ (e.g. $S^{eff}(k,p,\theta) = \sum_{i\in N} v^i(k,\theta) - p^i c$ for a designer interested in the maximization of social surplus).

Definition 6 *Given preferences S of the designer, we say that the feasible choice rule $(\psi,p) \in CR^f$* ***ex-ante dominates***[14] *the feasible choice rule (ψ',p') if*

$$E_\mu\left[S\left(\psi\left(\theta\right),p\left(\theta\right),\theta\right)\right] \geq E_\mu\left[S\left(\psi'\left(\theta\right),p'\left(\theta\right),\theta\right)\right]$$

We call $(\psi,p) \in CR^f$ ***ex-ante optimal*** *if no $(\psi',p') \in CR^f$ ex-ante dominates (ψ,p).*

In analogy to Holmstrom and Myerson (1983), we call $(\psi,p) \in CR^{act}$ (resp. $CR^{sim}, CR^{def}, CR^{seq}$) ***ex-ante incentive optimal*** *in the class of choice rules that are implementable with activating (resp. simultaneous, deferred, sequential) participation if no $(\psi',p') \in CR^{act}$ (resp. $CR^{sim}, CR^{def}, CR^{seq}$) ex-ante dominates (ψ,p).*

In mechanism design problems with $c = 0$, $CR^{act}, CR^{sim}, CR^{def}, CR^{seq}$ all coincide and the optimization over these sets reduces to an optimization over incentive compatible choice rules. If $c > 0$ we have seen that these sets of implementable choice rules generally do not coincide and that no set of direct mechanisms is equivalent to CR^{act} (or CR^{seq}).

This leaves us with three ways to identify an ex-ante incentive optimal mechanism in these classes:

Implementation of the first best: If we identify an ex-ante optimal $(\psi,p) \in CR^f$ and can show that $(\psi,p) \in CR^{act}$ we know a fortiori that (ψ,p) is ex-ante incentive optimal. This approach will allow us to identify the ex-ante efficient auction rule in the next section.

Implementation of second best: If we find the ex-ante incentive optimal choice rule (ψ,p) in CR^{def} and can show that actually $(\psi,p) \in CR^{act}$ we know a fortiori that (ψ,p) is ex-ante incentive optimal in CR^{act}. This approach will allow us to identify the ex-ante revenue-optimal auction rule in the next section.

Asymptotic characterization: In settings where the lower and upper bound (in CR^{sim} and CR^{def}) converge as $c \rightarrow 0$, the optimum in CR^{dir} will be "almost optimal" for small enough c.

3 Optimal auctions with participation costs

Consider a symmetric, private-values, one-object auction setting with partcipation costs c and the restriction that only participants in the auction game can receive the object and make payments[15], i.e.

- An agent's valuation of the object is equal to her private type: $v^i\left(\theta^i\right) = \theta^i \in \left[\underline{\theta},\overline{\theta}\right]$ for all $i \in N$

[14] Just like in Stegemann (1996) this seems to be a better optimality concept than ex-post optimality, i.e. there is no other feasible choice rule (ψ',p') with $S(\psi(\theta),p(\theta),\theta) \leq S(\psi'(\theta),p'(\theta),\theta)$ for all θ and strict inequality for a positive measure of θ.

[15] This assumption, phrased NRC by Stegemann (1996), is reasonable as we interpret the participation costs of the auction as the formal procedures that bidders have to undergo to prove their eligibility to receive the objects and make payments.

One could alternatively argue that only the participation in some bidding process is costly whereas, say, a random allocation could be effectuated without anybody incurring participation costs.

- Types are independently and identically distributed according to $d\mu(\theta) = \prod_i f^i(\theta^i)$ for smooth identical pdfs $f^i = f$ over $[\underline{\theta}, \overline{\theta}]$
- Virtual valuations $J(\theta^i) := \theta^i - \frac{1-F(\theta^i)}{f(\theta^i)}$
- $K = (N \cup \{s\}) \times \mathbb{R}^N$ with elements (a,t) where $a = s$ denotes the allocation that the seller keeps the object and $t = (t^i)_i$ are payments by the agents. Outcome assignment functions are accordingly noted as $\phi = (a,t) : T \to (N \cup \{s\}) \times \mathbb{R}^N$
- Utility is quasi-linear in money $u^i(k,\theta) = y^i v^i(\theta) - t^i$ where $y^i = y^i(k)$ indicates whether i receives the object, $y^i = 1$, or not, $y^i = 0$.
- Only participants can win the object: $K(I) \subseteq (I \cup \{s\}) \times (R^I \times 0^{N/I})$

This is the setting in which Stegemann (1996), Tan and Yilankaya (2006) study efficient auctions with simultaneous participation and Celik and Yilankaya (2007) study revenue-optimal auctions with simultaneous participation. These papers start from the observation that equilibrium strategies in the second price auction are characterized by cut-off values. If a bidder's valuation θ exceeds her cut-off value she participates and bids her valuation, otherwise she abstains. The surprising finding is that there may exist asymmetric equilbria, in which one bidder uses a high cut-off value and the other a low cut-off value. To see the reason for this consider first a symmetric equilibrium with two bidders, and then decrease the cut-off value for bidder i. This decreases the payoff from participation of the other bidder j who will thus respond by an increase in her cut-off value. This in turn increases the payoff from participation of i, supporting the participation of his weaker types. It turns out that the asymmetric equilibria of the second price auction often implement the ex-ante incentive optimal allocation rule (in CR^{sim}).

Ehrmann and Peters (1994) consider mechanisms with sequential participation in the above framework. The seller contacts the potential buyers in a random order and can stop the process at any time, and sell the object to a buyer he has contacted. They find that generally "fix price or best offer" with a "buy it now" price $\bar{p}$ and a reserve price $\underline{p}$ is an optimal selling mechanism. Upon being contacted, a potential buyer i can decide to either stay away from the mechanism and not incur the participation costs, or participate and buy the object immediately for $\bar{p}$, or submit a bid b^i in $[\underline{p}, \bar{p}]$. In the latter case she will win the object if no other bidder buys the object immediately for $\bar{p}$ or submits a bid higher than b^i. While a pure auction (setting $\bar{p} \geq \overline{\theta}$) is never optimal, the authors find distributions for which a pure fix price mechanism (setting $\bar{p} = \underline{p}$) is optimal. At first sight this seems odd, as one would expect the optimal price path to be decreasing because the option value of keeping the object decreases as the seller is contacting buyers unsuccessfully. Note, however, that such a mechanism with descending prices does not satisfy "sequential participation", as it assumes that the bidders can observe the current price when making their participation decision.

In the remainder of this section we will study optimal auction mechanisms in the larger classes of mechanisms with activating and deferred participation. Again, we would like to point out that we are interested in mechanisms with deferred participation not so much for their own sake but mostly as an upper bound for waht are implementable choice ruels with activating participation. We will see in the sequel that this upper bound can actually be attained, i.e. the optimal auction with activating participation achieves the same level of revenue (resp. social surplus) as the optimal auction with deferred participation.

We proceed as follows. In the following subsections we identify the direct revelation mechanisms with deferred participation that maximize social surplus and revenue. The results are straightforward and give us an upper bound for what is achievable with "real-life mechanisms". We next study which more reasonable auction formats, that satisfy activating participation, can achieve similar outcomes or even the same outcome.

In subsection 3.3 we model an open out-cry auction that satisfies activating participation and show that (a symmetric equilibrium in) this auction implements the optimal allocation rules identified above. Thus, a commonly observed auction format turns out to be optimal.

3.1 Social Surplus as optimality criterion

The ex-post efficient allocation rule in the auction setting assigns the object to the highest bidder, if her valuation, $\theta^{(1)}$, exceeds the cost of participation c. Consider the following direct revelation mechanism with deferred participation:

- Bidders report types $\widehat{\theta}^i$
- The auctioneer instructs the highest type $j = \arg\max_i \left\{\widehat{\theta}^i\right\}$ to participate, if $\widehat{\theta}^j \geq c$
- j receives the object for a price of $\theta^{(2)} - c$ if $\widehat{\theta}^j \geq c$, and nobody receives the object otherwise

Obviously we have:

Proposition 7 *The above mechanism is incentive compatible and implements the ex-ante efficient allocation rule.*

In the example of Stegemann (1996) with $N = 2$ and $\left[\underline{\theta}, \overline{\theta}\right] = [1, 2]$ this mechanism has an expected surplus of $5/3 - 0.3 = 1.37$. To the contrary, the optimal mechanism with simultaneous participation, i.e. the asymmetric equilibrium of the simultaneous second price auction, achieves a social surplus of 1.22.

3.2 Revenue as Optimality Criterion

To characterize the revenue-optimal auction we first adapt the revenue equivalence principle to our setting with participation costs.

Lemma 8 (Revenue Equivalence Principle) *Consider any equilibrium π of any mechanism m. Expected equilibrium transfers of any type $\overline{t}^i\left(\theta^i\right) := E_{\theta^{-i}}\left[t^i\left(\pi\left(\theta\right)\right)\right]$ are determined by the allocation function y and the participation function p. More specifically we have*

$$E_\theta\left[\sum_i t^i\left(\theta\right)\right] = \int_\Theta \sum_i \left(y^i\left(\theta\right) J^i\left(\theta^i\right) - p^i\left(\theta\right) c\right) d\mu - \sum_i \overline{u}^i\left(\underline{\theta}\right) \quad (1)$$

where $J^i\left(\theta^i\right) = \theta^i - \frac{1-F^i\left(\theta^i\right)}{f^i\left(\theta^i\right)}$ is the virtual valuation of agent i and $\overline{u}^i\left(\underline{\theta}\right)$ is the expected utility of type $\underline{\theta}$ of agent i in the mechanism.

Proof. Completely analogue to the common revenue equivalence principle. Fix π and m and denote by $\overline{u}^i\left(\theta^i\right)$ (resp. $\overline{y}^i\left(\theta^i\right), \overline{t}^i\left(\theta^i\right), \overline{p}^i\left(\theta^i\right)$) expected utility (resp. probability of receiving the object, payment, probability of entering the mechanism) given type θ^i

Consider the expected equilibrium utility of type θ^i:

$$\overline{u}^i\left(\theta^i\right) := E_{\theta^{-i}}\left[y\left(\pi^i\left(\theta^i\right), \pi^{-i}\left(\theta^{-i}\right)\right)v^i\left(\theta^i\right) - t^i\left(\pi^i\left(\theta^i\right), \pi^{-i}\left(\theta^{-i}\right)\right) - p^i\left(\pi^i\left(\theta^i\right), \pi^{-i}\left(\theta^{-i}\right)\right)c\right]$$

In equilbrium $\pi^i\left(\theta^i\right)$ maximizes this term, such that we get

$$\frac{\partial}{\partial\theta^i}\overline{u}^i\left(\theta^i\right) = E_{\theta^{-i}}\left[y\left(\pi^i\left(\theta^i\right), \pi^{-i}\left(\theta^{-i}\right)\right)\frac{\partial}{\partial\theta^i}v^i\left(\theta^i\right)\right] = \overline{y}^i\left(\theta^i\right)\frac{\partial}{\partial\theta^i}v^i\left(\theta^i\right)$$

Thus, we can write

$$\begin{aligned}\overline{t}^i\left(\theta^i\right) &= \overline{y}^i\left(\theta^i\right)v^i\left(\theta^i\right) - \overline{u}^i\left(\theta^i\right) - \overline{p}^i\left(\theta^i\right)c \\ &= \overline{y}^i\left(\theta^i\right)v^i\left(\theta^i\right) - \int_{\underline{\theta}}^{\theta^i}\overline{y}^i\left(\widetilde{\theta}^i\right)\frac{\partial}{\partial\theta^i}v^i\left(\widetilde{\theta}^i\right)d\widetilde{\theta}^i - \overline{u}^i\left(\underline{\theta}\right) - \overline{p}^i\left(\theta^i\right)c\end{aligned}$$

proving the first claim[16].

Applying $v^i\left(\theta^i\right) = \theta^i$, integration over θ^i, another application of independece of types, and a subsequent integration by parts yield the second claim. ■

The next step in the search for the revenue optimal mechanism is to maximize equation (1) through choice of $(y(\theta), p(\theta))$ subject to the constraints that:

1. $\overline{y}^i\left(\theta^i\right)$ is weakly increasing
2. $\sum_i y^i(\theta) \leq 1$ for all θ
3. If $p^i(\theta) = 0$ then also $y^i(\theta) = 0$, as by assumption only participants can receive the object
4. $y(\theta), p(\theta)$ arise from some equilibrium of some dynamic mechanism that satisfies the respective participation specification

As usual, we will for the moment ignore constraint 1 and also 4, maximize equation (1) pointwise subject to constraints 2 and 3 and then check that the ignored constraints are actually satisfied by the identified maximum.

First, we drop the term $-\sum_i \overline{u}^i\left(\underline{\theta}\right)$ as it is bounded above by zero by the interim participation constraints of the agents. Second, we realize that the constrained optimum of $\sum_i\left(y^i(\theta)J^i\left(\theta^i\right) - p^i(\theta)c\right)$ is achieved by $y^i(\theta) = p^i(\theta) = 1$ if $i \in \arg\max_j\left\{J^j\left(\theta^j\right) - c\right\} \geq 0$ and $y^i(\theta) = p^i(\theta) = 0$ else.

Consider the following direct revelation mechanism with deferred participation:

- Bidders report types $\widehat{\theta}^i$
- The auctioneer instructs the highest type $j = \arg\max_i\left\{\widehat{\theta}^i\right\}$ to participate, if $\widehat{\theta}^j \geq J^{-1}(c) - c$

[16] Note that up to here we only used quasi-linearity of utility function, convexity of the type space and independence of types, but not the 1-dimensionality of the type space or any property of the valuation function.

- j receives the object for a price of $\min\left\{r, \theta^{(2)} - c\right\}$, if $\widehat{\theta}^j \geq J^{-1}(c) - c$, and nobody receives the object otherwise

Obviously we have:

Proposition 9 *The above mechanism is incentive compatible and implements the ex-ante revenue optimal allocation rule.*

One should compare the performance of this mechanism with the optimal mechanisms with simultaneous participation, identified by Celik and Yilankaya (2007). For parameters $\theta \in [0,1], F(\theta) = \theta^4, c = 0.4$ they compute an optimal revenue of 0.2525 for the optimal asymmetric and 0.2515 for the optimal symmetric auction.

3.3 The Open Out-Cry Auction

We will now show how a very natural auction mechanism, that satisfies activating participation, implements the same outcome as the abstract direct mechanisms with deferred participation, identified in Propositions 7 and 9.

Consider the following model of a standard open out-cry auction with a reserve price r. The auction starts at time 0 and proceeds in real time t. At any time t there is a high price $p(t)$ (see below). At any time t a potential bidder i can choose to participate in the auction (at cost c) and to subsequently submit an intial bid $\underline{b}^i \geq p(t)$ and a maximum bid $\overline{b}^i \geq \underline{b}^i$. Once i has done so, she cannot take further actions in the auction at $t' > t$. The current price $p(t)$ is equal to 1) the reserve price r as long as no bidder has entered the auction, 2) the intial bid $\underline{b}^i$ if only bidder i has entered, and 3) the second highest maximum bid $\overline{b}^{(2)}$ if two or more bidders have entered the auction. This bidding procedure with an initial bid and a maximum bid is very similar to the one implemented on eBay.

The auction ends after a period Δ of inactivity at time $t + \Delta$ (where t denotes the time of the last entry) and awards the object the participant i with the highest $\overline{b}^i$ at the current price $p(t+\Delta)$. This participation structure satisfies our definition of "activating participation" as bidders only bid in the auction after they have incurred the participation costs. Before the costly participation decision they can only observe the behavior of their competitors.

We will now study an equilibrium of the open out-cry auction. Let $\beta(\theta)$ be the symmetric equilibrium bidding function in a first-price auction with reservation price r but without participation costs, and let $b(\theta) = \beta(\theta - c)$ be the same bidding function if valuations are decreased by c.

Consider now the following bidding strategy (b, t)

- If no bidder has bid at time $t\left(\theta^i\right) := \frac{\overline{\theta} - \theta^i}{\overline{\theta} - \underline{\theta}}\Delta$, bidder i submits the initial bid $\underline{b}^i = b\left(\theta^i\right)$ and maximal bid $\overline{b}^i = \theta^i$ at this time.
- If at time $t\left(\theta^i\right)$ the current high bid $p\left(t\left(\theta^i\right)\right)$ exceeds $b\left(\theta^i\right)$, bidder i does not enter the auction.
- If at time $t\left(\theta^i\right)$ the current high bid $p\left(t\left(\theta^i\right)\right)$ is less than $b\left(\theta^i\right)$, bidder i submits the initial bid $\underline{b}^i = b\left(\theta^i\right)$ and maximal bid $\overline{b}^i = \theta^i$ at this time.

Proposition 10 *The above strategies* $(b, t)_i$ *constitute a symmetric sequential equilibrium of the open out-cry auction. Thus, the open out-cry auction with* $r = 0$ *implements the ex-ante efficient allocation rule whereas the open out-cry auction with* $r = J^{-1}(c) - c$ *implements the revenue optimal allocation rule.*

It is interesting to note that this equilibrium of the ascending auction combines features of the equilibria of the 1st price auction, i.e. the shaded initial bid, and the sencon price auction, i.e. bidding up to one's true valuation after having entered. Note also that this equilibrium is not ex-post, i.e. it depends on correct beliefs of the agents. Thus the question around the optimal robust, i.e. ex-post incentive compatible, mechanism remains open.

The coordinated timing of participation, i.e. $t\left(\theta^i\right) := \frac{\overline{\theta}-\theta^i}{\overline{\theta}-\underline{\theta}}\Delta$, serves as an implicit Dutch price clock that shows price $b(\theta)$ at time $t(\theta)$. On the equilibrium path, bidder i who decides to enter at time $t\left(\theta^i\right)$ endogenously chooses the right bid $b\left(\theta^i\right)$. Although other bidders could overbid her they will not do so in equilibrium as they infer from i's entry that she has a higher valuation and would eventually win the auction anyway.

To prove the Proposition, consider for a moment a Dutch auction with an explicit price clock where the first bidder wins the auction.

Lemma 11 (Dutch Auction) $b(\theta) = \left(b^i\left(\theta^i\right)\right)_i$ *is a symmetric equilibrium in the Dutch auction.*

Proof. Note that in the Dutch auction, unlike in a sealed-bid first price auction, only the highest bidder actually submits her bid and thus only the highest bidder incurs the participation costs. Thus, given valuations θ^i and other bidders' bids b^{-i}, bidder i's maximization problem is

$$\max_{b^i \leq r}\left(\theta^i - b^i - c\right)\mathbb{I}_{\{b^i \geq b^{-i}\}} = \max_{b^i \leq r}\left(\left(\theta^i - c\right) - b^i\right)\mathbb{I}_{\{b^i \geq b^{-i}\}}$$

and thus identical to the maximization problem of the first price auction with participation costs, reduced valuations and the same reserve price. ■

Note that in principle we could declare the problem solved now, as the Dutch auction satisfies activating participation and implements the desired choice rules.[17]However, even though the Dutch auction formally satisfies the requirements of activating participation an explicit price clock is clearly an abstraction from reality in many auction settings, that given the real world prevalence of open out-cry auctions we think it worthwhile to show that this more applied mechanism also implements optimal allocation rules.

Proof of Proposition 10. The previous lemma shows there are no profitable deviations from (b, t) in which type θ^i mimicks type $\theta^{i\prime}$. We will now prove that there are no other profitable deviations either.

First, note that bidding one's true valuation as maximum bid $\overline{b}^i = \theta^i$ is always optimal, conditional on having participated by the usual argument in ascending auctions.

Next, consider the situation after some opponent j has entered before $t\left(\theta^i\right)$ with a bid $b\left(\theta^j\right) \geq b\left(\theta^i\right)$. From this entry, θ^i infers that $\theta^j \geq \theta^i$ and that thus entering the auction herself would be unprofitable. Note that this inference is valid both on and off the equilibrium path. On the other

[17] Moreover, the Dutch auction has the advantage of making the coordination of bidders explicit, rather than hoping that bidders will coordinate in equilibrium. However, note that even though entry coordination in the open out-cry auction will certainly not be perfect in reality, it is reasonable to assume that qualitatively strong bidders will enter early to pre-empt competition whereas weaker bidders will prefer to "wait and see" whether they have a chance of winning the auction, before investing resources into participation.

hand if some opponent j has entered before $t\left(\theta^i\right)$ with a bid $b\left(\theta^j\right) \leq b\left(\theta^i\right)$ off the equilibrium path, it is reasonable for i to assume that this is the action of a low type θ^j who hopes to "steal the auction" by entering early, i.e. this belief satisfies the conditions of sequential equilibrium and even forward induction equilibrium. Given these beliefs it is optimal for i to enter with bid $b\left(\theta^i\right)$ at time $t\left(\theta^i\right)$, anticipating to win the auction. Thus we can focus on deviations of θ^i that are not contingent on $-i$'s actions before $t\left(\theta^i\right)$.

Next consider deviations in which θ^i enters before $t\left(\theta^i\right)$, say at time $t\left(\theta^{i\prime}\right) < t\left(\theta^i\right)$. A bid of $\underline{b}^i = b\left(\theta^{i\prime}\right)$ at $t\left(\theta^{i\prime}\right)$ would mean mimicking type $\theta^{i\prime}$ and thus cannot constitute a profitable deviation by Lemma 11. A higher bid $\underline{b}^i \geq b\left(\theta^{i\prime}\right)$ would obviously be even less profitable. A lower bid, $\underline{b}^i = b\left(\theta^{i\prime\prime}\right) \leq b\left(\theta^{i\prime}\right)$, on the other hand could not be a best response to $(b,t)^{-i}$ for the following reason: This bid does not deter entry of types $\theta^j \in \left[\theta^{i\prime\prime}, \theta^{i\prime}\right]$. Therefore bidding $\underline{b}^i = b\left(\theta^{i\prime\prime}\right)$ at $t\left(\theta^{i\prime}\right)$ is dominated by waiting until $t\left(\theta^{i\prime\prime}\right)$ before entering at price $p\left(t\left(\theta^{i\prime\prime}\right)\right)$ and refraining from entering if $p\left(t\left(\theta^{i\prime\prime}\right)\right) \leq b\left(\theta^i - c\right)$.

Lastly, consider deviations in which θ^i enters after $t\left(\theta^i\right)$, say at time $t\left(\theta^{i\prime}\right) > t\left(\theta^i\right)$. The above arguments show that such a deviation can not be profitable if θ^i plans to stay away from the auction if the waiting strategy fails in the sense that an opponent j enters the auction in $\left[t\left(\theta\right), t\left(\theta^{i\prime}\right)\right]$. Thus we can focus on deviations in which θ^i plans to enter and win the auction if some type $\theta^j \in \left[\theta^{i\prime}, \theta^i\right]$ enters before her. The utility of θ^i of such a strategy is given by

$$u^i\left(\theta^i, \theta^{i\prime}\right) := \int_{\theta^{i\prime}}^{\theta^i} \left(\theta^i - c - \theta\right) dF^{N-1}\left(\theta\right) + \left(\theta^i - c - b\left(\theta^{i\prime}\right)\right) F^{N-1}\left(\theta^{i\prime}\right)$$

The net margin of winning against a competitor of type $\theta \in \left[\theta^{i\prime}, \theta^i\right]$ is $\theta^i - c - \theta$ whereas the net margin of winning when no opponent of type $\theta \geq \theta^{i\prime}$ has entered is given by $\theta^i - c - b\left(\theta^{i\prime}\right)$. Consider the marginal effect $-\frac{d}{d\theta^{i\prime}} u^i\left(\theta^i, \theta^{i\prime}\right)$ of decreasing $\theta^{i\prime}$ at $\theta^{i\prime} = \theta^i$. The marginal effect on $\int_{\theta^{i\prime}}^{\theta^i} \left(\theta^i - c - \theta\right) dF^{N-1}\left(\theta\right)$ is negative because the integrand is negative. The marginal effect on $\left(\theta^i - c - b\left(\theta^{i\prime}\right)\right) F^{N-1}\left(\theta^{i\prime}\right)$ is equal to zero as b is the equilibrium bidding function in the Dutch auction. Thus $u^i\left(\theta^i, \theta^i - \varepsilon\right) < u^i\left(\theta^i, \theta^i\right)$ for small values of ε. In other words, it is not worthwhile to wait ε longer at time $t\left(\theta^i\right)$ because entering the auction against an opponent $\theta^j \in \left[\theta^i - \varepsilon, \theta^i\right]$ is unprofitable as the margin of ε does not cover the participation costs of c.

Now consider $u^i\left(\theta^i, \theta^{i\prime} - \varepsilon\right) - u^i\left(\theta^i, \theta^{i\prime}\right)$ for any value of $\theta^{i\prime} \leq \theta^i$. It is easy to compute that

$$u^i\left(\theta^i, \theta^{i\prime} - \varepsilon\right) - u^i\left(\theta^i, \theta^{i\prime}\right) = \left(b\left(\theta^{i\prime}\right) - b\left(\theta^{i\prime} + \varepsilon\right)\right) F^{N-1}\left(\theta^{i\prime} + \varepsilon\right) - \int_{\theta^{i\prime} - \varepsilon}^{\theta^{i\prime}} \left(\theta - b\left(\theta^{i\prime}\right)\right) dF^{N-1}\left(\theta\right)$$

does not depend on θ^i as the allocation is not affected by this change of strategy, but only the price is. Therefore we get that $u^i\left(\theta^i, \theta^{i\prime} - \varepsilon\right) - u^i\left(\theta^i, \theta^{i\prime}\right) = u^i\left(\theta^{i\prime}, \theta^{i\prime} - \varepsilon\right) - u^i\left(\theta^{i\prime}, \theta^{i\prime}\right) < 0$ by the above. Thus $u^i\left(\theta^i, \theta^{i\prime}\right)$ is decreasing in $\theta^{i\prime}$, proving that a deviation of delaying entry to $t\left(\theta^{i\prime}\right) > t\left(\theta^i\right)$ and outbidding any opponent who enters in between is unprofitable as well. ■

4 Conclusion

We have introduced a dynamic mechanism design model with participation costs, applied it to a standard auction setting and found that an open out-cry auction with an appropriate reserve price is an optimal mechanism in this framework. In the presence of participation costs, dynamic mechanisms generally improve upon static mechanisms because they allow agents to take the costly participation decisions in a sequential, coordinated way. The open out-cry auction achieves this goal as in equilibrium only one bidder enters the auction and wins the object with her entry bid. The off-equilibrium entry threat of opposing bidders sustains a high bid of the winning bidder.

This leads us to suspect that the application to auctions, where only one bidder needs to enter, may not be representative for other multi-agent mechanism design settings such as voting setting or a bilateral trade setting. In these applications multiple agents need to participate to achieve desirable outcomes. The ensuing problems of free-riding (waiting for one's opponent to enter first) and signalling (entering in a bilateral trade setting is a signal of a high type) will likely make the analysis more difficult than in these applications.

References

[1] Bergemann, D. and Välimäki, J. (2002) "Information Acquisition and Efficient Mechanism Design", Econometrica, 70, 1007-1034

[2] Borgers, T. (2004) "Costly Voting," American Economic Review 94 (1), 57-66

[3] Bulow, J. and Klemperer, P. "When are Auctions Best", working paper

[4] Celik, G. and Yilankaya, O. (2007) "Optimal Auctions with Participation costs", discussion paper University of British Columbia

[5] Compte, and Jehiel, P. (2008) "Information acquisition in Auctions: Sealed-bid or Dynamic Formats?" forthcoming Rand Journal of Economics

[6] Cremer, J., Spiegel, Y. and Zheng Z. (2007) "Optimal Search Auctions", Journal of Economic Theory 134, 226-248.

[7] Cremer, J., Spiegel, Y. and Zheng Z. (2008) "Auctions with costly information acquisition", forthcoming in Economic Theory

[8] Ehrmann, C. and Peters, M. (1994) "Sequential Selling Mechanisms", Economic Theory 4, 237-253

[9] Gershkov, A. and Szentes B. (2008) "Optimal Voting Schemes with Costly Information Acquisition", forthcoming in Journal of Economic Theory

[10] Holmstrom, B. and Myerson, R. "Efficient and Durable Decision Rules with Incomplete Information", Econometrica 51, 1799-1819

[11] McAfee, P. and McMillan J. (1987) "Auctions with Entry", Economics Letters 23, 343-347

[12] McAfee, P. and McMillan J. (1988) "Search Mechanisms", Journal of Economic Theory 44, 99-123

[13] Myerson, R. (1981) "Optimal Auction Design", Mathematics of Operations Research 6, 58-73

[14] Samuelson, W. (1985) "Competitive bidding with entry costs", Economics Letters 17, 1-2, 53-57

[15] Stegemann, M. (1996) “Participation Costs and Efficient Auctions”, Journal of Economic Theory 71, 228-259

[16] Tan, G. and Yilankaya, O. (2006) "Equilibria in second price auctions with participation costs," Journal of Economic Theory 127, 205-219

[17] Zheng, C. “VCG Mechanisms in Dynamic Processes”, private communication