Report Number: WUCS-99-05

1999-01-01

Auctions without Common Knowledge

Sviatoslav B. Brainov and Tuomas W. Sandholm

This paper proves that the revenue equivalence theorem ceases to hold for auctions without common knowledge about the agents' prior beliefs. That is, different auction forms yield different expected revenue. To prove this, an auction game is converted to a Bayesian decision problem with an infinite hierarchy of beliefs. A general solution for such Bayesian decision problems is proposed. The solution is a generalization of the standard Bayesian solution and coincides with it for finite belief trees and for trees representing common knowledge. It is shown how the solution generalizes the frequently used technique of backward induction for infinite... Read complete abstract on page 2.

Auctions without Common Knowledge

Sviatoslav B. Brainov and Tuomas W. Sandholm

Complete Abstract:

This paper proves that the revenue equivalence theorem ceases to hold for auctions without common knowledge about the agents' prior beliefs. That is, different auction forms yield different expected revenue. To prove this, an auction game is converted to a Bayesian decision problem with an infinite hierarchy of beliefs. A general solution for such Bayesian decision problems is proposed. The solution is a generalization of the standard Bayesian solution and coincides with it for finite belief trees and for trees representing common knowledge. It is shown how the solution generalizes the frequently used technique of backward induction for infinite belief trees. The solution can be applied to any game with infinite belief trees. Computation of the solution does not rely on approximating the infinite trees with finite ones. The method can be used, for example, to analyze the expected revenue of alternative auction forms.

Auctions without Common Knowledge

Sviatoslav B. Brainov and Tuomas W. Sandholm

WUCS-99-05

January 1999

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Auctions without Common Knowledge

Sviatoslav B. Brainov and Tuomas W. Sandholm
Washington University
St. Louis, MO 63130
{brainov, sandholm}@cs.wustl.edu

Abstract

This paper proves that the revenue equivalence theorem ceases to hold for auctions without common knowledge about the agents' prior beliefs. That is, different auction forms yield different expected revenue. To prove this, an auction game is converted to a Bayesian decision problem with an infinite hierarchy of beliefs. A general solution for such Bayesian decision problems is proposed. The solution is a generalization of the standard Bayesian solution and coincides with it for finite belief trees and for trees representing common knowledge. It is shown how the solution generalizes the frequently used technique of backward induction for infinite belief trees. The solution can be applied to any game with infinite belief trees. Computation of the solution does not rely on approximating the infinite trees with finite ones. The method can be used, for example, to analyze the expected revenue of alternative auction forms.

1 Introduction

Auctions have been a subject of continuous interest in multiagent systems and electronic commerce [Monderer and Tennenholtz, 1998; Sandholm, 1996]. The allocative efficiency of auctions ensures their pervasive use in electronic markets. One of the main advantages of auctions as a form of market organization is their ability to cope with market imperfections. The most typical imperfections of electronic markets are the small number of market participants and the existence of incomplete and asymmetric information.

Usually in electronic commerce and multiagent systems, an auctioneer faces several possible buyers and has imperfect information about how much the buyers might be willing to pay. The problem of the optimal auction design [Myerson, 1981; Monderer and Tennenholtz, 1998] is to set up such auction rules that give the seller the highest possible utility.

Most theoretical results on optimal auction design draw crucially on the revenue equivalence theorem [Vickrey, 1961]. According to the theorem, the first-price sealed bid, second-price sealed bid, English and Dutch auctions are all optimal selling mechanisms provided that they are supplemented by an optimally set reserve price. The revenue equivalence theorem is based on the following assumptions: the bidders are risk neutral, payment is a function of bids alone, the auction is regarded in isolation of other auctions, the bidders' private valuations are independently and identically distributed random variables, every bidder knows only his own valuation and is uncertain about the other agents' valuations, there is common knowledge about the valuations' distribution. In this context common knowledge means that everybody knows the common prior distribution from where valuations are drawn, everybody knows that everybody knows, etc., ad infinitum.

The notion of common knowledge plays a central role in decision making, game theory and economics of uncertain information [Bacharach et al., 1997; Fagin et al., 1995; Geanakoplos, 1994; Brandenburger, 1993; Halpern and Moses, 1990]. While in game theory and many economic applications the common knowledge assumption might be applied innocuously, in electronic commerce there is no sufficient justification for the ubiquity of its use. A common feature of electronic commerce transactions is their anonymity. Only parties involved in the transaction have information about the transaction terms. Third parties are usually unable to access any specific transaction information and might even be unaware of the existence of the transaction. Therefore, tracking other agents' past behavior and forming expectations about their future behavior can be difficult. This argument can be extended to an extreme if we take into account the fact that autonomous agents usually act as intermediaries in electronic markets. It is possible that the same agent represents different parties at different moments of time. Therefore, in electronic markets where agents do not know about or cannot recognize one another, there are no sufficient grounds for applying the common knowledge assumption. It has also been formally shown that common knowledge is unobtainable by communication, no matter how much communication is allowed [Halpern and Moses, 1990].

In this paper the common knowledge assumption about prior beliefs is dropped, but all other classic assumptions are kept intact. In particular, the assumption that the agents' valuations are drawn from the same prior is kept. It is shown that without common knowledge the revenue equivalence theorem ceases to hold. The failure of revenue equivalence has significant practical importance since different auction forms lead to different expected revenues to the auctioneer.

The research presented in this paper is closely related to the work in game theory devoted to games with incomplete information [Harsanyi, 1967]. In our work the epistemic state of each agent is modeled as an infinite hierarchy of beliefs. Harsanyi [1967] suggested that each hierarchy of beliefs could be summarized by the notion of agent's type. Later Mertens and Zamir [1985] proved that the space of all possible types is closed in the sense that it is large enough to include even higher-order beliefs about itself. Brandenburger [1993] has shown that if agents' beliefs are coherent the space of all possible types is closed.

Our approach is related to the work of Gmytrasiewicz, Durfee and Vidal [Gmytrasiewicz and Durfee, 1995; Vidal and Durfee, 1996]. They presented a solution method based on finite hierarchies of beliefs. The main advantage of their recursive modeling method is that the optimal solution can always be derived. The recursive modeling method is based on the assumption that once an agent has run out of information his belief hierarchy can be cut at the point where there is no sufficient information. At the point of cutting, absence of information is represented with a uniform distribution over the space of all possible beliefs. The beliefs of order higher than the order of cutting are ignored. This approach, however, cannot be applied for rational agents with perfect reasoning abilities. We cannot prohibit such agents from forming higher-order beliefs by applying a uniform distribution whenever there is no sufficient information. Once an agent has run out of information at some level of beliefs, he has also run out of information for higher-order beliefs while continuing to model further the belief tree. Unlike the method of Gmytrasiewicz, Durfee and Vidal, our method allows such extended modeling by applying a decision-making procedure based on infinite hierarchies of beliefs, and leads to different results.

The paper is organized as follows. In Section 2 a simple auction setting is defined. The auction setting is used to exemplify the theoretical conclusions drawn in the later sections. In Section 3 a decision making model based on infinite hierarchies of beliefs is introduced. Analysis of auctions without common knowledge is presented in Section 4. Finally, the paper concludes by summarizing the results and providing directions for future research.

2 A Simple Auction Setting

In order to prove the failure of the revenue equivalence theorem, a simple auction setting is considered. The setting includes two risk-neutral bidders in an isolated auction for a single indivisible object. Suppose that each bidder has one of two possible valuations of the object: t_1 or t_2 (with $t_1<t_2$). Each bidder knows his own valuation, but is uncertain about his rival's valuation. Assume that valuations are independent and that there exists some objective distribution π from which valuations are drawn. Let π be common knowledge between bidders.

The setting so defined satisfies all the assumptions of the revenue equivalence theorem. Therefore, the first-price and the second-price sealed bid auctions yield for each bidder the same expected utility.

The assumption of common knowledge about prior beliefs does not affect the outcome of the second-price sealed bid auction. In that auction every bidder has a dominant strategy: bidding his own valuation. Bidding one's own valuation does not require anticipating the rival's behavior or holding any beliefs about the rival's beliefs.

On the other hand, the first-price sealed bid auction is sensitive to the common knowledge assumption. In such an auction, the agent's utility maximizing bid is a function of his beliefs about other agents' beliefs. The analysis of optimal bidding in such auctions is usually conducted using the Nash equilibrium solution concept from noncooperative game theory [Nash, 1951], or a refinement thereof. In such an equilibrium, each agent bids in a way that is a best response to the other agents' bidding strategies. However, the Nash equilibrium solution concept relies heavily on the common knowledge assumption. Up to now there has been no satisfactory equilibrium concept for games without common knowledge. One cannot derive the optimal bids in the first-price auction without such a solution concept. Therefore, one cannot calculate the expected utility of the bidders either. Thus, we need a solution concept for an auction game without common knowledge. Such a concept is proposed in the next section.

3 A Decision Making Model Based on an Infinite Hierarchy of Beliefs

In this section we propose a solution for a first-price sealed bid auction without common knowledge about bidders' prior beliefs. In Subsection 3.1 we discuss how bidders' prior beliefs can be represented by infinite hierarchies. Then we convert the simple auction game into a Bayesian decision problem based on an infinite hierarchy of beliefs. In Section 3.2 we propose a solution for the class of Bayesian decision problems based on infinite hierarchies of beliefs.

3.1 Infinite Hierarchies of Beliefs

Consider a first-price sealed bid auction without common knowledge about bidders' prior beliefs. Assume that all other strategically relevant information is common knowledge. This means that in our simple auction setting bidders have common knowledge about the two possible valuations t_1 and t_2, i.e. the support of the objective distribution π, but do not have common knowledge about π itself. Therefore, each bidder might hold some private beliefs about his rival's valuation distribution. For example, bidder i might believe that $\pi=(q,1-q)$. That is, bidder i might believe that bidder j's valuation equals t_1 with probability q and t_2 with probability 1-q. At the same time bidder j might believe that $\pi=(r,1-r)$, $r\neq q$. Meanwhile it might turn out that the actual distribution π differs from the bidders' beliefs.

The belief structure of each agent might be represented by a hierarchy of beliefs [Brandenburger, 1993]. Figure 1 is a tree diagram that represents the beliefs of an agent who knows that his valuation is t_1.

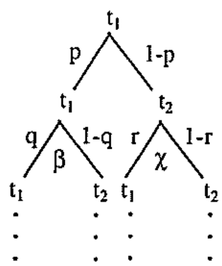

Figure 1: An infinite belief tree.

Suppose that the agent under consideration is agent i. First-order beliefs of agent i are represented by a discrete probability distribution σ_1, $\sigma_1=(p,1-p)$. Second-order beliefs are represented by two distributions. The first distribution σ_2, $\sigma_2=(q,1-q)$, corresponds to the case where agent j has valuation t_1. That is, agent i believes that if agent j's valuation is t_1, then agent j beliefs that $\pi=\sigma_2$. The second distribution σ_3, $\sigma_3=(r,1-r)$ is for the case where agent j's valuation is t_2. The belief hierarchy shown in Figure 1 is extended to infinity. Thus, every belief hierarchy over the space $T=\{t_1,t_2\}$ can be represented by a binary tree. The vertices of the tree are labeled with agents' valuations and the edges are labeled with probabilities. The valuation in the root of the tree corresponds to the valuation of the player whose beliefs are represented by the tree.

Let α, β, χ,... denote trees. The tree α represented in Figure 1 may also be represented in a list notation:

$$\alpha=t_1(p,\beta;1-p,\chi).$$

Here t_1 is the root of the tree, (p,1-p) is the probability distribution on the set of the immediate descendents of t_1 and β and χ are the subtrees whose roots are the immediate descendents of t_1.

Definition 1. The number of different probability distributions contained in a belief tree is called *belief power* of the tree.

The belief power of a tree can vary from 1 to infinity. The belief power of a belief tree often increases exponentially with the depth of the tree. It is useful to identify the class of belief trees for which the belief power is a linear function of the depth. In the following definitions the class of *k-uniform belief trees* is introduced. An infinite belief tree is k-order uniform if for every m, $1\leq m\leq k$, there is only one probability distribution at level m of the belief tree. In the following definitions $P^m(\alpha)$ denotes the unique probability distribution associated with level m of tree α.

Definition 2. Every belief tree α, $\alpha=t_i(p,\beta;1-p,\chi)$, is *first-order uniform*. By definition $P^1(\alpha)=(p,1-p)$.

Definition 3. A belief tree $\alpha=t_i(p,\beta;1-p,\chi)$, is *k-order uniform* if β and χ are k-1-order uniform and for every m, $1\leq m\leq k-1$, $P^m(\beta)=P^m(\chi)$. Then $P^n(\alpha)$ is defined as $P^n(\alpha)=P^{n-1}(\beta)$ for all n, $2\leq n\leq k$.

Definition 4. An infinite belief tree is *uniform* if it is k-order uniform for each $k\geq 1$.

Consider, for example, the tree represented in Figure 1. Since at the second level of the tree we have two probability distributions, $\sigma_2=(q,1-q)$ and $\sigma_3=(r,1-r)$, the belief tree is not uniform. In order to make it second-order uniform we have to set $\sigma_2=\sigma_3$.

The following propositions follow immediately from Definition 4.

Proposition 1. If the prior beliefs $\pi=(p,1-p)$ are common knowledge, then any infinite belief tree is uniform.[1]

Proposition 2. There is common knowledge about prior beliefs $\pi=(p,1-p)$ iff for every two belief trees α and β, such that β is a subtree of α, it holds that $P^1(\alpha)=P^1(\beta)$.

Now we are in a position to convert our simple auction game into a Bayesian decision problem based on an infinite hierarchy of beliefs.

Definition 5. *A Bayesian decision problem* for agent i is given by: (i) the set $T=\{t_1,t_2\}$, the possible valuations of the opponent; (ii) S_i, a compact set of all strategies available to agent i; (iii) $U_i : S_i x T \rightarrow R$, utility function of agent i; and (iv) α, an infinite belief tree of agent i.

In the next subsection we propose a solution to a Bayesian decision problem based on an infinite hierarchy of beliefs.

3.2 Solution to a Bayesian Decision Problem based on an Infinite Belief Hierarchy

Most of the research in Bayesian decision theory is based on finite belief trees. Bayesian decision problems based on infinite belief hierarchies are studied by Tan and Verlang [1988] and Armbruster and Boge [1979]. In order to cope with the infinite recursion of beliefs these studies impose the so called minimum consistency requirement. According to this requirement if the probability of an event is computed using k levels of the belief tree or m levels, they must give the same result.

In this paper we propose a solution which does not rely on the minimum consistency requirement. The solu-

[1] Due to space limitations the more straightforward proofs are omitted in this version of the paper.

tion is a generalization of the solution of Tan and Werlang and can be applied to finite as well as to infinite belief trees. The solution coincides with the standard Bayesian solution for finite trees and for trees representing common knowledge.

Let T be the set of all bidder valuations. In our example $T=\{t_1,t_2\}$. The class of all infinite belief trees over T is denoted by G. Every α, $\alpha \in G$, is represented as a pair $(V(\alpha),E(\alpha))$, where $V(\alpha)$ is the set of vertices and $E(\alpha)$ is the set of edges. Let $r(\alpha)$ denote the root vertex of α. Recall that vertices are labeled with valuations and edges are labeled with probabilities. For a belief tree α, $\alpha \in G$, we denote the vertex labeling function by n_α, $n_\alpha: V(\alpha) \rightarrow T$.

Let S be the strategy set of an agent. With each vertex of the belief tree we assign a strategy which tells what the decision maker would do at that vertex if he were there. Formally we denote the strategy labeling by $\phi_\alpha: V(\alpha) \rightarrow S$ and for every vertex v, $v \in V(\alpha)$, $\phi_\alpha(v) \in S$.

For each infinite belief tree there exists an infinite number of strategy labelings. However, only few of them (if any) meet the Bayesian rationality requirement, i.e., that each strategy has to be a best response to the profile of all other strategies. In the following definition the class of all strategy labelings is restricted to the class of balanced strategy labelings. They satisfy the Bayesian rationality requirement. A strategy labeling ϕ_α is balanced if the strategy associated with each vertex is a best response to the strategies associated with the successor vertices, given the probabilities assigned to the successors. Formally,

Definition 6. A strategy labeling of α, $\alpha \in G$, is *balanced* iff for each subtree $\beta=(p,\chi;1\text{-}p,\delta)$ of α (including α) it holds that $\phi_\alpha(r(\beta))$ is a best response to the mixture of strategies $[p, \phi_\alpha(r(\chi)); 1\text{-}p, \phi_\alpha(r(\delta))]$.

Definition 6 provides a solution concept for a Bayesian decision problem based on an infinite belief hierarchy. For finite belief trees this concept coincides with the standard Bayesian solution. The concept of balanced strategy labeling preserves the central principle of consistency in the sense of Hammond [1988]. The central principle of consistency says that the decision maker's decision at a vertex in a tree should depend only on the part of the tree that originates at that vertex. The central principle of consistency justifies the frequently used technique of backward (bottom-up) induction (recursion). The concept of balanced strategy labeling generalizes the backward induction to the case of infinite trees. If we have derived a strategy labeling for some level of a tree we can "cut" the belief hierarchy at that level and apply backward (bottom-up) induction starting from the cutting level. By doing so we do not lose any strategically relevant information, since the concept of balanced labeling guarantees that the strategies along the cutting line convey all the relevant information belonging to the infinite part of the tree.

The following proposition applies immediately to our simple auction example. It provides necessary and sufficient conditions for the existence of equilibrium in an auction with two possible bidder valuations.

Proposition 3. Suppose that the prior beliefs $\pi=(p,1\text{-}p)$ are common knowledge. A balanced strategy labeling exists iff there are two strategies $s_1 \in S_{t1}$ and $s_2 \in S_{t2}$ such that:
s_1 is a best response to the strategy mixture$[p,s_1;1\text{-}p,s_2]$
s_2 is a best response to the strategy mixture $[p,s_1;1\text{-}p,s_2]$.

When the prior beliefs are not common knowledge, the following definition can be useful for finding a balanced strategy labeling.

Definition 7. A uniform infinite belief tree α allows common knowledge from level k, $k \geq 1$, iff $P^k(\alpha)=P^{k+1}(\alpha)=\ldots$

According to Definition 7, a belief tree allows common knowledge from level k if all belief subtrees which start at level k imply common knowledge. That is, after some nesting of beliefs the bidder "begins" to believe that there is common knowledge about priors.

The following procedure uses Definition 7 to find a balanced strategy labeling. Suppose that the tree α, $\alpha \in G$, allows common knowledge from a given level k, $k>1$. Then we may "cut" α at level k and apply Proposition 3 to all infinite subtrees of α starting at level k. By doing so we find a balanced strategy labeling for the levels greater or equal to k. After that, since the remaining part of the tree is finite, we may apply backward (i.e., bottom-up) induction starting at the level k-1 and ending at the root of the tree.

4 Application of the Decision Making Model to Auction Analysis

In this section we analyze the first-price sealed bid auction without common knowledge about prior beliefs. We restrict our analysis to the auction setting defined in Section 2. All assumptions made in Section 2 hold.

4.1 The Case with Common Knowledge

Before proceeding to the case without common knowledge, we look for a solution for the first-price auction where there is common knowledge about prior beliefs. The solution is provided by the following proposition.

Proposition 5. Suppose that the prior beliefs p and 1-p are common knowledge. Then for the first-price sealed bid auction the bidder's expected utility is 0 when the bidder's valuation is t_1 and $p(t_2\text{-}t_1)$ when the bidder's valuation is t_2.

Proof. Since there does not exist an equilibrium in pure strategies, we look for an equilibrium where each bidder with valuation t_1 bids t_1 ($t_1<t_2$), and each bidder with valuation t_2 randomizes according to a continuous cumulative distribution function F(x) with continuous support on $[a_1,a_2]$, where $t_1 \leq a_1 \leq a_2 \leq t_2$. It can be shown that this equilibrium is unique. Clearly, $a_1=t_1$. If $a_1>t_1$, then a bidder with valuation t_2 would be better off bidding $t_1+\varepsilon$ rather than bidding a_1. In order for a bidder with valuation t_2 to play a mixed strategy in the interval $[a_1,a_2]$ he must be indifferent *ex ante* between all bids in this interval. Hence, for every bid $x\in[a_1,a_2]$ it holds that

$$(t_2-x)(p+(1-p)F(x))=c,$$

where c is constant. Here t_2-x is the bidder's utility if he wins and $p+(1-p)F(x)$ is the probability of winning. Because $F(t_1)=0$, it follows that $c=(t_2-t_1)p$. Thus, the continuous distribution function $F(x)$ is implicitly defined by

$$(t_2-x)(p+(1-p)F(x))=(t_2-t_1)p \tag{1}$$

Substituting a_2 for x in Equation (1) and taking into account that $F(a_2)=1$, we obtain

$$a_2=pt_1+(1-p)t_2.$$

Therefore, the bidder's expected utility equals 0 when his valuation is t_1 and $(t_2-t_1)p$ when his valuation is t_2. □

4.2 The Case without Common Knowledge

Suppose now that there is no common knowledge about prior beliefs. Each bidder holds some private first-order beliefs about t_1 and t_2. Suppose further that the bidders have no additional information about one another. This is a realistic assumption, since in many electronic commerce applications bidders cannot identify one another. The absence of information might be represented for example as a uniform distribution over the set T [Gmytrasiewicz and Durfee, 1995].[2] That is, a bidder who is not inclined to believe that one of the outcomes t_1 or t_2 is more likely may tend to assign equal probability to both outcomes. In that case, the bidder's second-order beliefs can be represented by a uniform distribution. Since the bidders are rational, we cannot prevent them from forming higher-order beliefs. The basic assumption for forming higher-order beliefs is the following: once a bidder has run out of information at level k, he also runs out of information at all levels m, m>k. According to this assumption, all higher-order beliefs are also represented by uniform distributions. A generic belief tree for a bidder with valuation t_1 is shown in Figure 5.

The solution for the case without common knowledge about prior beliefs is provided by the following theorem. Surprisingly, the first-order beliefs about priors do not affect the optimal bidding strategy.

Proposition 6. When the prior beliefs are not common knowledge and bidders run out of knowledge for second-order beliefs, the first-price sealed bid auction yields expected utility 0 to the bidder with valuation t_1 and $\frac{1}{2}(t_2-t_1)$ to the bidder with valuation t_2. The optimal bid and the expected utility do not depend on the bidders' first-order prior beliefs.

[2] In general, our solution concept does not rely on such a uniformity assumption.

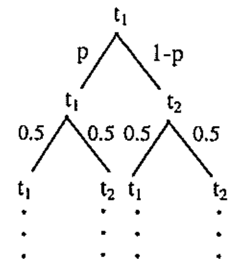

Figure 5: A generic belief tree.

Proof. Analogous to the case with common knowledge we look for equilibrium where every bidder with valuation t_1 bids t_1 and every bidder with valuation t_2 bids some value not less than t_1. Therefore a bid equal to t_1 might be assigned to all t_1-vertices of the belief tree. Since the strategies of the bidders with valuation t_1 are known, all infinite trees starting at t_1-vertices can be cut off. The resulting tree is shown in Figure 6.

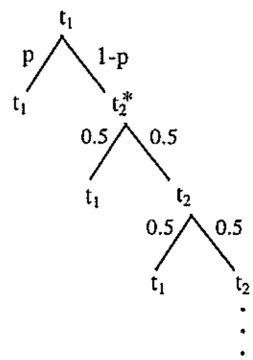

Figure 6: The resulting belief tree.

Now we are interested in finding a balanced strategy labeling for the tree in Figure 6. Since we have already labeled t_1-vertices, it remains to obtain a labeling for all t_2-vertices. Consider the infinite subtree whose root is marked with asterisk in Figure 6. This tree allows common knowledge from level 1 and therefore we can apply Propositions 3 and 5. Thus, for this subtree there exists a balanced strategy labeling. This labeling assigns bid t_1 to each t_1-vertex and randomized strategy in the interval $[t_1, \frac{1}{2}(t_1+t_2)]$ to each t_2-vertex. The distribution function $F(x)$ of the randomized strategy is defined by the following equation

$$(t_2-x)(\tfrac{1}{2}+\tfrac{1}{2}F(x))=\tfrac{1}{2}(t_2-t_1) \tag{2}$$

What remains to be done is to find a bidding strategy b* corresponding to the root of the tree. It is clear that b* must be a best response to the strategy mixture $[p,t_1;1-p,b^{**}]$, where b** is the strategy defined by Equation (2). We can solve Equation (2) for $F(x)$, thereby obtaining

$$F(x)=(x-t)/(t_2-x).$$

The expected utility of submitting bid x, given that the rival adheres to the strategy mixture $[p,t_1;1-p,b^{**}]$ is:

$(t_2-x)(p+(1-p)(x-t_1)/(t_2-x))$ when $t_1 \leq x \leq ½(t_1+t_2)$ or
t_2-x when $½(t_1+t_2)<x$.

In order to obtain an optimal bid we have to maximize the expected utility function. There are three possible cases:

(i) $p<½$: the optimal bid is $½(t_1+t_2)$. The expected utility is $½(t_2-t_1)$;

(ii) $p=½$: every bid in the interval $[t_1,½(t_1+t_2)]$ is optimal. The expected utility is $½(t_2-t_1)$;

(iii) $p>½$: the optimal bid is $½(t_1+t_2)$. The expected utility is $½(t_2-t_1)$. ❒

Theorem 1. When there does not exist common knowledge about private beliefs, the revenue equivalence theorem ceases to hold, i.e., the bidder's expected utility is different for different types of auctions.

Proof. To prove the failure of the revenue equivalence theorem, it is sufficient to find two auctions which give the bidders different expected utility. Consider the first-price sealed bid auction and the second-price sealed bid auction. It follows from Proposition 6 that the expected utility for the bidder with valuation t_2 is $½(t_2-t_1)$ in the first-price sealed bid auction without common knowledge about prior beliefs. On the other hand, for the second-price auction the optimal strategy for every bidder is to bid his own valuation. Therefore, in the second-price auction the expected utility for the bidder with valuation t_2 is $(t_2-t_1)p$, where p is subjective probability that the other bidder's valuation is t_1. Thus, when $p\neq½$, the two auctions yield different expected utility. ❒

5 Conclusions

In this paper a solution for a Bayesian decision problem based on an infinite belief hierarchy was presented. The solution is a generalization of the standard Bayesian solution and coincides with it for finite belief trees and for trees representing common knowledge. The computation of our solution does not rely on approximating the infinite belief trees by finite belief trees.

It was shown that without common knowledge about prior beliefs the fundamental revenue equivalence theorem ceases to hold. The failure of the revenue equivalence theorem has significant practical importance. Since different auctions yield different revenues, auction designers should be careful when choosing auction rules. This opens promising prospects for comparative analysis of different auction forms using the solution concept presented in this paper.

References

[Armbruster and Boge, 1979] W. Armbruster and W. Boge. Bayesian game theory. In O. Moeschlin and D. Pallaschke, editors. *Game Theory and Related Topics.* North-Holland, Amsterdam, 1979.

[Bacharach et al., 1997] Michael Bacharach, Louis-Andre Gerard-Varet, Philippe Mongin and Hyun Shin. *Epistemec Logic and the Theory of Games and Decisions.* Kluwer, 1997.

[Brandenburger, 1993] Adam Brandenburger. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59:189--198, 1993.

[Fagin et al., 1995] Ronald Fagin, Joseph Halpern, Yoram Moses, Moshe Vardi. *Reasoning About Knowledge.* MIT Press, Cambridge, Massachusetts, 1985.

[Halpern and Moses, 1990] Joseph Halpern and Yoram Moses. Knowledge and common knowledge in distributed environment. *Journal of ACM*, 37(3):549--587,1990.

[Hammond, 1988] P. J. Hammond. Consequentialist foundation for expected utility. *Theory and Decision*, 25:25--78, 1988.

[Harsanyi, 1967] John Harsanyi. Games with incomplete information played by Bayesian players, Part I. *Management Science*, 14(3):159--182, 1967.

[Geanakoplos, 1994] John Geanakoplos. Common knowledge. In K. J. Arrow and M.D. Intriligator, editors. *Handbook of Game Theory with Economic Applications.* Elsevier, 1994.

[Gmytrasiewicz and Durfee, 1995] Piotr J. Gmytrasiewicz and Edmund H. Durfee. A rigorous, operational; formalization of recursive modeling. *Proceedings of the First International Conference on Multi-Agent Systems*, pages 125--132, 1995.

[Mertens and Zamir, 1985] Jean-Francois Mertens and Shmuel Zamir. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory,* 14:1--29, 1985.

[Monderer and Tennenholtz, 1998] Dov Monderer and Moshe Tennenholtz. Optimal auctions revisited. *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, pages 32--37, 1998.

[Myerson, 1981] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58--73, 1981.

[Nash, 1951] John F. Nash. Noncooperative games. *Annals of Mathematics*, 54:28--295, 1951.

[Sandholm, 1985] Tuomas W. Sandholm. Limitations of the Vickrey auction in computational multiagent systems. *Proceedings of the Second International Joint Conference on Multi-Agent Systems*, pages 229--306, 1996.

[Tan and Werlang, 1988] Tommy Tan and Sergio Werlang. The Bayesian foundations of solution concepts of games. *Journal of Economic Theory*, 45:370--391, 1988.

[Vickrey, 1961] William Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8--37, 1961.

[Vidal and Durfee, 1996] Jose Vidal and Edmund Durfee. The impact of nested agent models in an information economy. *Proceedings of the Second International Conference on Multi-Agent Systems*, pages 377--384, 1996.