

Impersonation Strategies in Auctions

Ian A. Kash and David C. Parkes

CRCS and SEAS, Harvard University

`kash@seas.harvard.edu`, `parkes@eecs.harvard.edu`

Abstract. A common approach to analyzing repeated auctions, such as sponsored search auctions, is to treat them as complete information games, because it is assumed that, over time, players learn each other's types. This overlooks the possibility that players may impersonate another type. Many standard auctions (including generalized second price auctions and core-selecting auctions), as well as the Kelly mechanism, have profitable impersonations. We define a notion of impersonation-proofness for the auction mechanism coupled with a process by which players learn about each other's type, and show an equivalence to a problem of dominant-strategy mechanism design.

Keywords: Auctions, Impersonation, Kelly Mechanism, Ad Auctions.

1 Introduction

Most analyses of auctions emphasize uncertainty. While a bidder may know his value for an item, he is unlikely to know exactly how every other bidder values it. However, he is likely to have some beliefs about others, and the standard Bayesian analysis of auctions requires that, in equilibrium, bidders act optimally based on their beliefs about other bidders.

This approach is natural for a single, stand-alone auction. However, in some cases (for example in sponsored search auctions) the same bidders will participate in many auctions. Thus, a notion of equilibrium should take into account that, over the course of many auctions, bidders will learn about each others valuations. Unfortunately, as the folk theorem shows, the set of potential equilibria in such a repeated setting is large and complicated.

One natural class of equilibria are those where players spend some time learning until they reach an equilibrium of the “stage game,” after which they use the same strategies forever. If players are no longer learning, then it seems reasonable that they have complete information about the types of other players.

However, this analysis glosses over a key point. These complete information equilibria will only be reached if players *correctly* learn each other's types. As the learning process is part of the repeated game, players may have an incentive to deviate during this process. This could be prevented by finding learning algorithms that are themselves an equilibrium of the repeated game. This is the approach taken by, for example, Brafman and Tennenholtz [3] and Ashlagi et

al. [2]. However, such algorithms require that most or all of the players participate and that they learn in a particular fashion, so it seems unlikely that they will be a good predictor of real-world behavior.

Our approach, in the same spirit as complete information analysis, is to assume that players will learn and reach an equilibrium. In particular, we ignore the possibility that players will do something other than learn the types of other players. We also ignore their rewards during the learning period and assume they only care about the long-run behavior that the complete information game naturally captures. With these assumptions, should we expect to reach a complete information equilibrium? In this paper, we argue that the answer is no. In particular, players have the option to *impersonate* another type and participate in the learning algorithm as if their true type were the type they are impersonating. This causes the other players to believe they are playing a different complete information game and so a different strategy profile is reached.

Complete information equilibrium analysis has been used for many auctions, notably Kelly [7] for bandwidth allocation; Edelman, Ostrovsky and Schwartz [5] for generalized second price auctions; and Day and Milgrom [4] for core-selecting auctions. All turn out to have profitable impersonation strategies. We define a notion of *impersonation-proofness* and show that it is equivalent to selecting equilibria that implement the outcome of a dominant strategy mechanism for the incomplete information problem.

2 Model

Consider a Bayesian game G . Each of n players i has a type $\theta_i \in \Theta_i$ drawn according to the joint distribution $F(\theta_1, \dots, \theta_n)$, which is common knowledge. Each player chooses an action $a_i \in A_i$ based on his type. Each player's utility, which may depend on the joint action and his type, is $u_i(a, \theta_i)$. A Bayesian Nash equilibrium is then defined in terms of an expectation over the types of players.

If players play this game repeatedly, we expect them to learn about their opponents. Theorems have been established regarding the Nash equilibria of the complete information game G_θ for a number of different Bayesian games G , where G_θ is G with $\theta = (\theta_1, \dots, \theta_n)$ made common knowledge.

We model this learning process as a *mediator*: players submit a type and the mediator suggests an action for each player. To keep in mind our intuition of players learning, we require that if players report θ to the mediator, the mediator suggests a Nash equilibrium of G_θ , as a goal of most learning dynamics is to reach equilibrium¹ [8]. Formally, given a Bayesian game G , a mediator is a function $M : \Theta \rightarrow A$ such that $M(\theta)$ is a Nash equilibrium of G_θ , where $\Theta = \Theta_1 \times \dots \times \Theta_n$ and $A = A_1 \times \dots \times A_n$.

Given a Bayesian game G and a mediator M , we have the mediated game G_M . First, each player i learns θ_i and submits some θ'_i to M . Then i learns $M_i(\theta')$ and selects an action a_i . This formulation suggests the obvious strategy

¹ Our formalism of a mediator is inspired by that of Ashlagi et. al [1], but our motivation and definition are slightly different.

of lying to the mediator in the first stage. We call such strategies *impersonation strategies* because in practice they amount to impersonating some other type for a period of time to convince other players that the player is actually of that type. We focus on impersonation strategies in a strong sense: the player not only lies to the mediator but then follows the mediator's advice based on that lie. Thus, the player can continue this impersonation indefinitely. A player has a *profitable impersonation* when he can increase his payoff by using an impersonation strategy when all other players report truthfully and follow the mediator's advice. Formally, i has a profitable impersonation if there exists some θ'_i such that

$$u_i(M_i(\theta'_i, \theta_{-i}), \theta_i) > u_i(M_i(\theta), \theta_i). \quad (1)$$

With this in mind, we say a mediated game G_M is *impersonation-proof* if no player ever has a profitable impersonation. Formally, for all i , θ , and θ'_i ,

$$u_i(M_i(\theta), \theta_i) \geq u_i(M_i(\theta'_i, \theta_{-i}), \theta_i). \quad (2)$$

3 Example: The Kelly Mechanism

As mentioned in the introduction, many games have profitable impersonations. In this short paper, we analyze one such example. Suppose the owner of a network wants to allocate bandwidth to users of the network. Kelly [7] introduced a simple mechanism for this problem. Each player i submits a bid b_i . He then receives a $b_i / \sum_j b_j$ fraction of the bandwidth and pays a cost of b_i . This mechanism has the nice property that each player needs only submit a bid rather than describe his entire, potentially complicated, utility function. Furthermore, if all players have concave utility functions, then there is a unique complete information Nash equilibrium which can be found using a simple learning algorithm². Johari and Tsitsiklis [6] showed that this mechanism has a price of anarchy of $4/3$.

The following lemma (whose proof is omitted) shows that it is quite common for players to have profitable impersonations. In particular, this means that, despite having a good price of anarchy, actual performance could be poor.

Lemma 1. *Consider the Kelly mechanism with two players who have linear utility functions ($u_i(x_i) = \theta_i x_i$) with $\theta_1, \theta_2 > 0$. Unless $\theta_1 = \theta_2$, both players have a profitable impersonation.*

To illustrate Lemma 1, suppose $\theta_1 = 2$ and $\theta_2 = 1$. Then the unique equilibrium has bids $(4/9, 2/9)$ so player 1's utility is $8/9$. Now suppose player 1 impersonates $\theta_1 = 3$. Now the unique "equilibrium" has bids $(9/16, 3/16)$ and player 1's utility is $15/16$. Thus player 1 has gained by pretending to have a higher valuation.

² There are different models of how players optimize for this mechanism. We assume players are *price-anticipating*: they take into account how their bid affects the price they pay when determining their optimal bid.

4 Impersonation-Proofness

In this section, we examine when mediated games are impersonation-proof and thus it is plausible that players would be willing to participate as their true type. Consider a Bayesian game G in which each player i chooses an action a_i and then his utility is determined by the vector of actions a and his type θ_i . For example, in a first price auction an action is a bid and the vector of bids determines the winner and each player's payment. In problems of interest, G is induced as the result of a designed mechanism, with a set of outcomes O , a set of joint actions A , a mapping $o : A \rightarrow O$, and utility functions $u_i(a, \theta_i) = u_i(o(a), \theta_i)$. We refer to $S = (\Theta, O, u)$ as the social choice problem domain.

Any mediator M for G is a function from type vectors to action vectors, and thus when combined with the mapping $o : A \rightarrow O$ is itself a mechanism. In fact, this is a *direct revelation* mechanism. A mediator coupled with a game defines a subset of the space of possible direct revelation mechanisms, insisting that $M(\theta)$ be a complete information Nash equilibrium for all θ .

Theorem 1. *Let G be a Bayesian game that is a mechanism (not necessarily incentive compatible) for a social choice problem domain $S = (\Theta, O, u)$. There exists an impersonation-proof mediator M for G iff there exists a dominant strategy mechanism D for S such that for all θ there exists an $a(\theta)$ that is a Nash equilibrium for G_θ and $D(\theta) = o(a(\theta))$*

Theorem 1 suggests a general approach to finding impersonation-proof mediators: take a dominant strategy mechanism D for the same problem domain and find equilibria that implement $D(\theta)$ in each game G_θ . For example, the bidder-optimal locally envy-free equilibrium of a generalized second price auction implements the VCG outcome [5], so the mediator that selects this equilibrium is impersonation-proof.

References

1. Ashlagi, I., Monderer, D., Tennenholtz, M.: Mediators in position auctions. *Games and Economic Behavior* 67(1), 2–21 (2009)
2. Ashlagi, I., Tennenholtz, M., Monderer, D.: Robust learning equilibrium. In: *Proc. of the 22nd Conference in Uncertainty in Artificial Intelligence, UAI* (2006)
3. Brafman, R.I., Tennenholtz, M.: Efficient learning equilibrium. *Artif. Intell.* 159(1–2), 27–47 (2004)
4. Day, R., Milgom, P.: Core-selecting package auctions. *Int. J. Game Theory* 36, 393–407 (2008)
5. Edelman, B., Ostrovsky, M., Schwartz, M.: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review* 97(1), 242–259 (2007)
6. Johari, R., Tsitsiklis, J.N.: Efficiency loss in a network resource allocation game. *Math. Oper. Res.* 29(3), 407–435 (2004)
7. Kelly, F.P.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
8. Shoham, Y., Powers, R., Grenager, T.: If multi-agent learning is the answer, what is the question? *Artificial Intelligence* 171(7), 365–377 (2007)