Games with incomplete information

Maryam Kamgarpour

May 21, 2025

1 Outline

- Review of auctions
- Nash equilibria in auctions
- Bayes Nash equilibrium
- Revenue equivalence for first-price and second-price auctions

2 Review of first-price and second-price auctions

Le us start by reviewing the first-price auction and the VCG mechanism. We will consider spectrum auctions, commonly employed by governments for allocating the right to use certain frequency bands to communication companies. You can read about Switzerland's spectrum auction Here.

Exercise 1. Spectrum Auction

Consider an auction to sell frequency spectrum licenses for 10 and 20 MHz in the East and West directions. The four possible options are shown below.

	Geographic direction	
Bandwidth	West-20	East-20
	West-10	East-10

There are five bidders — these can be telecommunication providers. In our problem, each bidder is interested for a particular combination of these options. Let their bids be denoted by x_i, $i = 1, \dots, 5$ and be as follows in million CHF:

$$\begin{aligned}
x^1(\text{West-20, East-20}) &= 90, \\
x^2(\text{West-10, East-10}) &= 100, \\
x^3(\text{West-20, East-20}) &= 110, \\
x^4(\text{East-20, East-10}) &= 100, \\
x^5(\text{West-20, West-10}) &= 100.
\end{aligned}$$

The interpretation of the above is that bidder 1 desires to have the 20 MHz spectrum in both the West and East directions and is willing to pay 90 mCHF for it. Determine the winners and the price that each winner pays to the auctioneer in

1. the pay-as-bid mechanism (first-price auction)

2. the VCG mechanism

Solution. In both mechanisms, the items are allocated to generate the highest profit. Let $\alpha^i \in \{0,1\}$ denote the decision variable corresponding to the bid of player i getting accepted. The feasible allocations are:

- $\alpha^1 = \alpha^2 = 1, \;\; \alpha^3 = \alpha^4 = \alpha^5 = 0$, which gives $J(\delta) = 190$
- $\alpha^2 = \alpha^3 = 1, \;\; \alpha^1 = \alpha^4 = \alpha^5 = 0$, which gives $J(\delta) = 210$
- $\alpha^4 = \alpha^5 = 1, \;\; \alpha^1 = \alpha^2 = \alpha^3 = 0$, which gives $J(\delta) = 200$

Thus, the optimal solution is $\alpha^2 = \alpha^3 = 1, \;\; \alpha^1 = \alpha^4 = \alpha^5 = 0$.

1. In the pay-as-bid mechanism, bidders 1, 4, and 5 will pay nothing. Bidders 2 and 3 will make a payment of 100 and 110 mCHF, respectively.
2. In the VCG mechanism, Bidders 1, 4, and 5 pay nothing. Bidder 2 pays $p_2 = 200 - (210 - 100) = 90$ mCHF. The analogous formula for bidder 3 gives $p_3 = 200 - (210 - 110) = 100$ mCHF. The total utility the auctioneer makes is 190 mCHF.

Note that our auction setting is arguably simple - the real auctions are much more complex. The point of our simple example is to help understand the main concepts. Usually, once we get the fundamentals, we can build more complexities on solid grounds. We saw that in a second-price auction, the dominant strategy Nash equilibrium is to bid truthfully. What about in a first-price auction? What is a Nash equilibrium bidding strategy?

Exercise 2. Let $t^1 > t^2 > \cdots > t^N$ denote the true valuations of bidders $1, 2, \ldots, N$ for an item being auctioned. Verify that in a first-price auction, a Nash equilibrium strategy is given by $x^1 = t^2 + \epsilon$, where ϵ is a small positive number (related to bid increments allowed) and $x^j = t^j$ for $i \neq 1$. In words, the bidder with the highest valuation of the item should bid the second highest price and other bidders should bid truthfully.

Solution. We can verify that no player has incentive to unilaterally deviate from the above strategies. Clearly, for any $x^1 > t^2$ player 1 still wins the auction but has to pay higher amount so she has a lower utility. For $x^1 < t^2$ player 1 no longer wins the auction. For all other players, for any $x^j > t^j$ the players could potentially win the auction but will have negative utility. However, for $x^j \leq t^j$, they will have zero utility.

What does "true" valuation mean? We saw that if the bidders in an auction are power producers, this true valuation can be their production marginal cost. In a spectrum auction, the companies may estimate their true valuation of a given frequency band by computing their potential benefits of using the band for broadcasting their information, for example, the profit generated from the advertisements on their channel. In a painting auction, the true valuation of each bidder depends on who they are. For example, a national gallery might have certain value for a piece by a famous artist, and a millionaire aiming to complete her/his art collection might have a different value for the art piece. In any case, we assume that each bidder has some true value for the item.

Obviously, the players do not know others' valuation of the items. They have *incomplete information* about the game. Hence, they cannot compute the Nash equilibrium strategy above. One may then only consider auctions such as second-price or VCG whose dominant strategy Nash equilibrium is truthful bidding. However, such mechanisms in general suffer from other shortcomings.

Alternatively, we can try to analyze games with incomplete information. We will shortly introduce this class of games and return to the question of computing equilibria for a first-price auction.

We will consider the Bayesian approach to analyze games with incomplete information. In this approach, all the unknowns in the model are captured by an uncertainty, and a prior distribution on the uncertainty set is assumed to be common knowledge. The original formulation of Bayesian games is due to [21]. Ever since, the topic has become standard in most game theory courses and books [22, 13, 17].

3 Bayesian games

Bayesian games were introduced by John Harsanyi to capture games with incomplete information. He won the Nobel Prize in Economics in 1994 and here you can read his Nobel Prize lecture on games with incomplete information.

Definition 1 (Bayesian game)**.** A Bayesian game consists of N players.

- A set $K = K^1 \times \cdots \times K^N$ of action spaces, with K^j being the action space for player j.
- A set of types $T = T^1 \times \cdots \times T^N$. Player j's type is $t^j \in T^j$ and is known only to her.
- A prior probability distribution D on the set of types T.
- A set of utilities $(J^1, \ldots, J^N)$, where $J^j : T \times K \to \mathbb{R}$ is player j's utility.

In the Bayesian formulation above, we assume the unknowns in the game are captured by the "type" of players. And all players have the same common prior distribution D about the types (this can be generalized to consider players having different priors on the types).

In a Bayesian game, each player must choose its action $x^j \in K^j$ not knowing others' types t^{-j}, but only the distribution D over all types. How should she select x^j? The idea is to associate an action for each type of the player. We define a pure strategy $s^j : T^j \to K^j$ as a map from player j's type to its action space. Given that player j knows her type t^j, she can use Bayes' rule[1] to compute the probability of others' types conditioned on her type, namely, $D^j := D(t^{-j}|t^j)$, so D^j is the distribution of types of other players, conditioned that type of player j is t_j. Her expected utility then is given by

$$E_{D^j} J^j(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})) := \sum_{t^{-j}} D^j(t^{-j}|t^j) J^j\big(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})\big).$$

Note 1: there is a slight abuse of notation in using $s^{-j}(t^{-j})$ because each player's strategy should be a function of her own type.

Note 2: Wile we assume the set of types T has finite cardinality. The generalization to uncountable types would simply replace the sum above with an integral.

We will use the expected utility above to define a Bayesian Nash equilibrium strategy:

Definition 2 (Bayesian Nash equilibrium)**.** A strategy profile $\{s^j\}_{j=1}^N$ with $s^j : T^j \to K^j$ is a Bayesian Nash equilibrium if each s^j is a best-response strategy to s^{-j} for all possible types t^j, that is,

$$E_{D^j} J^j(t^j, t^{-j}, s^j(t^j), s^{-j}(t^{-j})) \geq E_{D^j} J^j(t^j, t^{-j}, x^j, s^{-j}(t^{-j})), \tag{1}$$

for all $x^j \in K^j$, $t^j \in T^j$, and for all $j \in \{1, \ldots, N\}$.

[1] recall that for two events A, B, Bayes' rule gives probability of event A conditioned on event B as follows: $P(A|B) = \frac{P(A \cap B)}{P(B)}$

		player 2	
		cooperate (C)	defect (D)
player 1	cooperate (C)	2, 2	0, 3
	defect (D)	3, 0	1, 1

Figure 1: Payoff matrix if player 2 is selfish

		player 2	
		cooperate (C)	defect (D)
player 1	cooperate (C)	3, 3	1, 2
	defect (D)	2, 1	0, 0

Figure 2: Payoff matrix if player 2 is nice

Example 1. Let us consider a simple two-player game, with a payoff matrix (utilities rather than losses) given as in Figure 1. In this table, the green colors correspond to the utility of the first player and the blue colors correspond to the utility of the second player. The first player chooses the row of the matrix and the second player chooses the column of the matrix. Furthermore, both players are maximizing their utility. Verify that (defect, defect) is the dominant strategy Nash equilibrium.

Now consider the situation where player 2 has two types. She can be either selfish or nice. In the selfish case, the payoff matrix is as per Figure 1. In the nice case, the payoff matrix is as per Figure 2. Player 1 does not know the type of player 2. How shall we define the concept of Nash equilibrium? One way to handle this uncertainty is to assume a probability for each type of player 2. We can say that with probability d player 2 is selfish and with probability $1-d$ player 2 is nice. Verify that in the Bayes Nash equilibrium, player 2 chooses D if selfish and C if nice, whereas player 1 will choose C if $d \leq 1/2$.

Solution. First, note that $T^2 = \{\text{selfish}, \text{nice}\}$ and player 1 does not have types. Hence, we look for a strategy for player 2 for each of her types and an action for player 1. What is the dominant strategy of player 2 if she is selfish? what about the case when she is nice?

First, observe that if player 2 is selfish, she has a dominant strategy to defect (D) and if she is nice, her dominant strategy is to cooperate (C). Note that this dominant strategy is $s^2 : T^2 \to \{C, D\}$. In other words, it is a strategy for each of player 2's type. Player 2 knows her type and accordingly will play her dominant strategy. Player 1 on the other hand has only one type. Hence, we need to determine his optimal action (not strategy since he has one type). She also does not know the type of player 2, only the probability of her being nice or selfish. The expected utility of player 1 from cooperating is $d \times 0 + (1-d) \times 3 = 3 - 3d$, and her expected utility from defecting is $d \times 1 + (1-d) \times 2 = 2 - d$. Hence, player 1 should cooperate for $d \leq \frac{1}{2}$ and should defect otherwise.

Note that player 1 does not have a type. Thus, for any prior distribution on player 2 being selfish (captured by d above), player 1 has an equilibrium action, rather than a strategy.

In general, computing Bayesian Nash equilibrium is extremely difficult. However, under certain assumptions and for certain classes of games, this computation can be done as we will see below for auctions.

4 First and second-price auctions revisited

Our goal now is to use the Bayesian game formulation to determine the Bayesian Nash equilibrium in the first-price auction as follows.

Exercise 3. Consider a first-price auction with two players, call them Alice and Bob. Alice's true valuation of the item is $a \in [0, 1]$ and Bob's true valuation of the item is $b \in [0, 1]$. We assume these true valuations are independently and uniformly distributed on the interval $[0, 1]$. Verify that $s^1(a) = \frac{a}{2}$ and $s^2(b) = \frac{b}{2}$ is a Bayesian Nash equilibrium bidding strategy for Alice and Bob, respectively. What is the Bayesian Nash equilibrium in this setting under the second-price auction?

Solution. Without loss of generality (due to symmetry), let us compute the best response strategy of Alice considering that Bob bids $b/2$. Let x be Alice's bid. The utility of Alice is $J^a = a - x$ if she wins the bid and $J^a = 0$ otherwise. Hence, the expected utility of Alice is

$$E_{D^a} J^a(a, b, x, \frac{b}{2}) = E_{D^a}\{(a - x)\mathbf{1}_{\{z \,|\, z > \frac{b}{2}\}}(x)\},$$

where $\mathbf{1}_C(x)$ is the indicator function of the set C. Since x and a are known to Alice, the only probabilistic variable is b, the type of Bob. Hence, expected utility of Alice is $(a-x) \times E_{D^a}\{\mathbf{1}_{z>\frac{b}{2}}(x)\}$. Now, since a and b are independent, $D^a = D(b|a)$ becomes the probability density of b. Hence, expected utility of Alice becomes equivalent to $(a - x) \times$ (Probability that $x > \frac{b}{2}$). In words, $(a - x) \times$ Probability that Alice wins). This was a lengthy discussion for a statement that may seem intuitive. Nevertheless, the point was to practice using the framework introduced, since the framework will become handy for more complex setups.

Let's compute probability that Alice wins. This is the probability that $b/2 < x$. Considering that b is uniformly distributed on $[0, 1]$, this probability is given as follows:

$$P(b < 2x) = \begin{cases} 0 & \text{if } x \leq 0, \\ 2x & \text{if } 0 \leq x \leq 1/2, \\ 1 & \text{if } x \geq 1/2. \end{cases}$$

Since utility of Alice is $a - x$ if she wins, it follows that she must compare the maximum of $a - x$ for $x \geq 1/2$ or the maximum $(a - x)2x$ subject to $0 \leq x \leq 1/2$. The first maximization gives the utility of $a - 1/2$ for $x = 1/2$ and the second one gives $a^2/2$ for $x = a/2$. As $a^2/2 - (a - 1/2) \geq 0$, it follows that the optimal solution is given by $x = a/2$. Hence, Alice's best response to Bob's bid of $b/2$ is to bid $a/2$. By symmetry, the same argument holds for Bob. Hence, $(a/2, b/2)$ is the Bayesian Nash equilibrium strategy.

For the second-price auction, we already saw that bidding truthfully is dominant strategy Nash equilibrium for each player. The fact that this is a dominant strategy, means that each player has no incentive to deviate from truthful bidding regardless of types (and hence corresponding actions) of the other players. Hence, Alice would bid a, while Bob would bid b.

The above analysis can be generalized to the first-price auction with N players, where each player's valuation is independently and identically distributed. The analysis becomes a little more involved, but there is a closed form solution for the uniform distribution.

From the auctioneer perspective, would the first-price or the second-price auction generate more profit? In our simple two-player example with uniform distribution of valuations on $[0, 1]$, in the case of first-price auction, the auctioneer's profit is $\max\{a/2, b/2\}$, with a and b uniformly distributed. In the second price auction, since the Nash equilibrium strategy is bidding truthfully, then the auctioneers revenue is $\min\{a, b\}$ (since the highest bidder gives the second high price).

Exercise 4. Verify that the expected profit of the auctioneer in both first-price and second-price auction under the two player setup above is $1/3$.

Solution. Hint: Start by computing the expected value of $c = \max\{a/2, b/2\}$, which is the auctioneer's profit in first-price auction, and $c = \min\{a, b\}$, the auctioneer profit in second-price auction, for a, b being uniformly distributed in $[0, 1]$.
Let us start with the first-price function. Let $c = \max\{a/2, b/2\}$. Consider the cumulative distribution function of the random variable c, denoted by $F(c \leq y)$. Now, $c \leq y \iff a \leq 2y \wedge b \leq 2y$, and since a and b are independent, $F(c \leq y) = P(a \leq 2y)P(b \leq 2y) = (2y)(2y) = 4y^2$. Hence, the probability density function of this random variable is $8y$. It follows that expected value of c is $\int_0^{1/2} y(8y)dy = 8y^3/3\big|_0^{1/2} = 1/3$. Similarly, you can compute the probability density function of the random variable $c = \min\{a, b\}$. Here, it might be easier to compute $F(c \geq y) = 1 - F(c \leq y)$, since $c \geq y \iff a \geq y \wedge b \geq y$. Then, verify that expected value of this random variable is $1/3$.

The above result is an instance of the celebrated *Revenue Equivalence Principle*. This theorem shows that for all mechanisms satisfying certain assumptions the expected utility of the auctioneer at a Bayesian Nash equilibrium is the same. For more details, see Chapter 9 of [13] and the paper of Myerson [23] — Myerson received the 2007 Nobel Prize in Economics, Nobel Prize lecture here[2]. Note that in practice, players may not have independent and identically distributed valuations. However, analyzing equilibria in auctions in more general settings is a very challenging problem. This problem has been receiving increasing attention from the computer science community due to their applications in online ad auctions.

5 Summary and further reading

There are several topics we have not had a chance to discuss in this course. In the references, we suggest additional text books for game theory as well as references for auctions. We will briefly comment on two of the most relevant research topics.

5.1 Learning Nash equilibria

Related to the incomplete information setting, a very active research topic is how players learn Nash equilibria in iterative games. Note that if players cannot learn to play Nash equilibria, then any theoretical analysis of outcome of the game based on the Nash equilibrium concept is questionable. In an auction such as those in electricity markets or for adverts on internet, players have generally less information than that assumed by a Bayesian game. In particular, they do not know how many other players there are, their strategy spaces, and sometimes they don't even fully understand the auction being run, specially in the case when there are multiple items or indivisible items and complex constraints (recall electricity market procurement constraints). So, players cannot compute their Bayesian Nash equilibrium strategy. Nevertheless, these auctions are repeated over and over. In each run, players are gathering information about the game based on their past observations. How should they use these observations and update their strategies?

This is the topic of learning Nash equilibria with limited information. In general, first one needs to assume the specific feedback players receive after each iteration of the game. Then, to devise a learning rule and prove that this rule converges to an equilibrium. A good reference is [6]. Under assumption of convex games with monotone pseudo-gradients (certainly limiting in terms of

[2] How Many Nobel Laureates we have met in this course?

auctions), I have been working on this topic in [7]. Without the above assumption, we also have some progress, see for example, [8]. For more general games with dynamic population, you can read [25].

5.2 Mechanism design

From the auctioneer or system operator's perspective, how should the market be designed to maximize social welfare? To do so, we need to incentivize bidders to bid truthfully. It is known that the VCG mechanisms are the only individually rational dominant strategy incentive compatible and efficient mechanisms. We have already discussed the issue with collusion and shill bidding. Given that it is theoretically impossible to ensure all desired properties of a mechanism, we might need to relax certain properties of the mechanism, such as incentive compatibility. The analysis of modified mechanisms becomes extremely challenging. It is in general better to resort to simulation to understand equilibrium strategies and the auction revenue/loss [28].

References

[1] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.

[2] J. B. Rosen, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica: Journal of the Econometric Society, 1965.

[3] S. Kakutani, *A generalization of Brouwer's fixed point theorem*, Duke mathematical journal, 8(3), 457-459, 1941

[4] F. Facchinei and J. Pang, *Finite-dimensional variational inequalities and complementarity problems*, Springer Science & Business Media, 2007.

[5] D. Paccagnan, M. Kamgarpour and J. Lygeros, *On Aggregative and Mean Field Games with Applications to Electricity Markets*, Proceedings of the European Control Conference, 2016.

[6] N. Cesa-Bianchi, G. Lugosi, *Prediction, learning, and games*, Cambridge university press; 2006 Mar 13.

[7] T. Tatarenko and M. Kamgarpour, *Payoff-Based Approach to Learning Generalized Nash Equilibria in Convex Games*, in IEEE Transactions on Control Systems, available at `https://arxiv.org/abs/1703.04113`, , 64(4), 1426-1439, 2018.

[8] P.G. Sessa, I. Bogunovic, A. Krause, M. Kamgarpour, *Contextual games: Multi-agent learning with side information*, Advances in Neural Information Processing Systems. 2020;33:21912-22.

[9] B. Gentile, F. Parise, D. Paccagnan, M. Kamgarpour and J. Lygeros, *Nash and Wardrop equilibria in aggregative games with coupling constraints*, to appear in IEEE Transactions of Automatic Control Problem, `https://arxiv.org/abs/1702.08789`, 2018.

[10] D. Monderer and S. Shapley. *Potential games*, Games and economic behavior 14.1 (1996): 124-143.

[11] G. Scutari, S. Barbarossa, and D. P. Palomar, *Potential games: A framework for vector power control problems with coupled constraints*, IEEE International Conference on Acoustics Speech and Signal Processing, volume 4, pages 241-244, May 2006.

[12] D. Bertsimas, Dimitris, V. Gupta, and I. Paschalidis, *Data-driven estimation in equilibrium using inverse optimization*, Mathematical Programming 153, no. 2 (2015): 595-633.

[13] N. Nisan, T. Roughgarden, E. Tardos, V. V. Vazirani, eds., 2007. *Algorithmic game theory.* Cambridge university press.

[14] L. M. Ausubel and P. Milgrom, *The lovely but lonely Vickrey auction*, Combinatorial auctions, 17, 22-26, 2006.

[15] P. Cramton, *Spectrum auction design*, Review of Industrial Organization, 42(2), 161-190, 2013.

[16] D. Easley and J. Kleinberg, *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*, `http://www.cs.cornell.edu/home/kleinber/networks-book/`, 2010.

[17] Y. Shoham and K. Leyton-Brown, *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*, Cambridge University Press, 2008.

[18] O. Karaca, P. G. Sessa, N. Walton, and M. Kamgarpour, *Designing coalition-proof reverse auctions over continuous goods*, `https://arxiv.org/abs/1711.06774`, 2017.

[19] E. Wolfstetter, *The Swiss UMTS spectrum auction flop: Bad luck or bad design?*, 2001.

[20] P. Klemperer, *How (not) to run auctions: The European 3G telecom auctions*, European Economic Review, 46(4-5), 829-845.

[21] J. C. Harsanyi, *Games with incomplete information played by "Bayesian" players, I-III Part I. The basic model*, Management science. 1967 14(3):159-82.

[22] M. J. Osborne and A. Rubinstein, *A course in game theory*, MIT press; 1994 Jul 12.

[23] R. B. Myerson, *Mechanism design*, in Allocation, Information and Markets 1989, pp. 191-206, Palgrave Macmillan, London.

[24] R. B. Myerson and M. A. Satterthwaite, *Efficient mechanisms for bilateral trading*, Journal of economic theory, 29(2):265-81, 1983.

[25] T. Lykouris, V. Syrgkanis, E. Tardos, *Learning and efficiency in games with dynamic population*, in Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms, pp. 120-129, Society for Industrial and Applied Mathematics, 2016.

[26] A. Dreves, F. Facchinei, C. Kanzow and S. Sagratella, *On the solution of the KKT conditions of generalized Nash equilibrium problems*, SIAM Journal on Optimization, 21(3):1082-108, 2011.

[27] F. Facchinei and C. Kanzow, *Generalized Nash equilibrium problems*, 4OR. 5(3):173-210, 2007.

[28] B. Bünz, B. Lubin and S. Seuken, *Designing core-selecting payment rules: A computational search approach*, `http://www.ifi.uzh.ch/ce/publications/Core_Payment_Rules.pdf`, 2018.