

Answer all questions. You have 3 hours in which to complete the exam. Don't spend too much time on any one question.

1. (30 Minutes – 20 Points)

Consider an asymmetric Cournot model. There are N firms with marginal costs $c_1 \leq c_2 < \dots \leq c_N$. The market inverse demand curve is $P(Q) = 1 - Q$.

(a) Assuming that the costs are such that all firms do make positive sales in equilibrium, solve for the equilibrium quantities q_1^*, \dots, q_N^* and the equilibrium price. What more primitive assumption on costs could have been used to ensure that the firms all produced in equilibrium?

(If you can't do this, just solve the model in the special case when the N firms have a common cost c . I'll give you partial credit on this part and it is enough to let you give complete answers to (b) and (c).)

(b) Suppose the costs are common across firms: $c_1 = c_2 = \dots = c_N = c$. Find the equilibrium profits as a function of N . How many firms would choose to enter in a pure strategy equilibrium if the firms needed to pay an entry cost of E to get into the market?

(c) Give a brief informal discussion of whether entry tends to be too low, socially efficient, or excessive in imperfect competition models with entry costs. Without doing any calculations what can you say about entry in this model relative the social optimum?

(d) Suppose that after all N firms have entered (all with marginal cost c), firm 1 gets an opportunity to invest $I^2/2$ to reduce its marginal cost to $c - kI$. Suppose that the investment level firm 1 chooses is observed by the other firms before they choose their quantities. What level of investment does firm 1 choose? Without doing any calculations, discuss whether firm 1 would invest more or less if its investment was not observed by the other firms? Without doing any calculations, discuss whether firm 1 would invest more or less if the investment opportunity arose before the firms made their entry decisions.

2. (30 Minutes – 20 Points)

Consider an asymmetric private value auction. Player 1's value is uniformly distributed on $[0, 2]$ and player 2's value is uniformly distributed on $[1, 3]$.

(a) What does the revenue equivalence theorem say about whether a pure first-price auction, a pure second-price auction, and a first-price auction with a reserve price raise the same revenue for a seller given these value distributions?

(b) Suppose the seller decides to use a second price auction with a reserve price, i.e. the object is sold at the maximum of r and the second highest bid if the winning bidder is willing to pay this, and not sold otherwise. What are the bidders' equilibrium strategies? Argue that a reserve price of less than one is never optimal. What is the seller's revenue when the reserve price r is between one and three?

(c) Suppose the seller also has a second unit to sell. It is common knowledge that player 1 would get no incremental benefit from winning a second unit. Player 2 would get the same benefit from the second unit that he gets from the first (what this is remains private information and uniformly distributed on $[1, 3]$.) Discuss informally why a simultaneous multiunit ascending auction (i.e. having people raise two hands if they want 2 units at p and one hand if they want one unit at p) can be problematic in a situation like this. In what senses is it a bad auction design?

(d) If you can, find the equilibrium of the auction described in part (c). What is the probability that the goods are allocated efficiently?

(Hint: player 2's first order condition is equivalent to being indifferent to first order between dropping out at $b^*(v)$ and $b^*(v) + db$.)

3. (30 Minutes – 20 Points)

In this question you need to make up the problem as well as to solve it.

The model(s) you make up are supposed to illustrate some basic principles of patent design: some patent protection is necessary to allow firms to invent; one doesn't want to grant excessive patent protection because this creates too much deadweight loss; and the government can use patent breadth as well as duration to achieve its goals.

As a starting point, assume that if a firm invests E it will invent a new product. Assume that there is a heterogeneous population of potential consumers. A type θ consumer gets utility $v - p + \theta$ if he buys the product at price p .

You can specify things like v , the distribution of θ , and what competitors can do under various patent protection regimes in any way you like to answer the questions.

(a) Write down a simple model in which patents provide complete protection against competition for some time period T . Show that the firm invests if and only if T is above some cutoff.

(b) Does your model have the property that government maximizes social welfare by making the patent duration as short as possible subject to the constraint that the firm be willing to invest? If not, how could you modify the model to get it to bring out the basic investment incentives vs. deadweight loss intuition.

(c) Now introduce some notion of patent breadth, x . The model should be such that when x is small competitors can produce very similar products and when x is large competitors' products must be less similar. Does your model have the property that the optimal policy for the government is to set $T = \infty$ and to make x as small as possible subject to the constraint that the firm be willing to invest? If not, discuss why it does not and how you might modify the model to make it have this property.

4. (40 Minutes – 20 points)

A number of recent papers have described how to estimate dynamic models where the equilibrium is Markov-perfect. A dynamic model in this setting consists of a state vector (s), a set of rules describing the transitions between states, and payoffs accruing to each player at each state (π). There are a finite number of infinitely-lived agents who discount the future at a constant rate (β), and time is discrete.

1. Bajari, Benkard, and Levin (BBL) suggest a general two-step estimation scheme for recovering the economic primitives in such a setting. In two sentences or less, explain what the researcher does (in the broadest sense) in step one and in step two.
2. Why do we break the process into two steps? How is BBL's approach an extension of what Guerre, Perrigne, and Vuong did in the auction setting?

Suppose we are interested in investigating the dynamic competition between Coke and Pepsi. Further, we observe the number of television commercials, A , in all markets across several decades. There are some time periods where neither firm advertises. Market j varies only by population, N_j . Prices are fixed at one for both products in all markets and marginal costs of production are zero. From in-depth research into the cola market, you know that quantities sold by Coke in a given period are represented by:

$$q_{Coke} = N_j \left(\frac{\exp(s_{Coke})}{1 + \exp(s_{Coke}) + \exp(s_{Pepsi})} \right),$$

where s_i is the state variable representing that firm's build-up stock of brand loyalty. The equations giving sales for Pepsi and the outside good are symmetric. No other state variables or trends enter the demand for cola products.

The profit function for each firm is

$$\pi_i = q_i - c(A),$$

where $c(A)$ is the unknown cost function of advertising that is of fundamental interest.

3. Explain why you can identify and recover the unobservable state variables using Berry's method without appealing to a dynamic model. Why does the outside good's state have to be normalized to zero for identification?
4. Assuming that advertising today only affects state variables tomorrow, explain how you would recover the transition matrix between states as a function of the number of current television commercials. Can you do this nonparametrically?
5. Suppose you know that your model is correctly specified, there is no error in the measurement of the data, and yet you still find error in the previous step. How can you rationalize this from the perspective of the underlying model?
6. After recovering this transition matrix, you know: (i) the states of both firms at all times in all markets, (ii) how much a television commercial shifts demand in the next

period, and (iii) the associated payoffs as a function of the state. This leaves one empirical component (iv) to recover in the first stage. What is it, and how would you recover it?

7. Explain why empirical objects (i)-(iv) are sufficient to recover the unknown cost function $c(A)$ in step two.
8. Intuitively, how are the costs of advertising, $c(A)$, nonparametrically identified?
9. If firms have differential levels of brand loyalty decay, what moments in the data identify those different rates?

5. (40 Minutes – 20 points)

Guerre, Perrigne, and Vuong (2000), hereafter GPV, demonstrated an efficient and non-parametric method for recovering valuations from the outcomes of auctions. Denote the utility of winning by U , the set of bids as B , and the set of bidders by N . Each bidder i , after observing signal $X_i = x_i$, solves the following optimization problem:

$$\max_{\tilde{b}} \left(E \left[U_i | X_i = x_i, \max_{j \in N_{-i}} B_j \leq \tilde{b} \right] - \tilde{b} \right) Pr \left(\max_{j \in N_{-i}} B_j \leq \tilde{b} | X_i = x_i \right). \quad (1)$$

Define bidder i 's expectation of its valuation conditional on its signal and the highest competing bid as:

$$\tilde{v}_i(x_i, m_i; N) = E \left[U_i | X_i = x_i, \max_{j \in N_{-i}} B_j = m_i \right]. \quad (2)$$

The expectation is conditional on the signal and on being pivotal, since equilibrium bidders will bid such that they expect to just win (ignoring ties, as they are probability zero events). Denote the conditional distribution of maximal opposing bids by $G_{M_i|B_i}(m_i|b_i; N)$, with associated density $g_{M_i|B_i}(m_i|b_i; N)$. We can then rewrite the optimization problem as:

$$\max_{\tilde{b}} \int_{-\infty}^{\tilde{b}} \left[\tilde{v}_i(x_i, m_i; N) - \tilde{b} \right] g_{M_i|B_i}(m_i|b_i; N) dm_i. \quad (3)$$

Leibniz's rule for differentiating integrals is:

$$\frac{\partial}{\partial z} \int_{a(z)}^{b(z)} f(x, z) dx = \int_{a(z)}^{b(z)} \frac{\partial f}{\partial z} dx + f(b(z), z) \frac{\partial b}{\partial z} - f(a(z), z) \frac{\partial a}{\partial z}.$$

1. Intuitively, what two things does the bidder optimally trade off when solving Equation 3?
2. Using Leibniz's rule, calculate the FOC for the bidder in Equation 3 and solve for the optimal bid: \tilde{b} .
3. What data do you need from a set of ex-ante identical auctions to identify the set of underlying valuations in (i) the symmetric affiliated private value model and (ii) the asymmetric affiliated private value model? Briefly, explain why.
4. The first step in GPV requires estimating $G_{M_i|B_i}(m_i|b_i; N)$ and $g_{M_i|B_i}(m_i|b_i; N)$. However, they actually suggest using two related distributions. What are they? Why can we substitute them? Why would we want to?
5. In the special case of symmetric (potentially affiliated) private values, we drop individual subscripts and note that only the number of bidders, n , and not their identities matter. One estimating equation is then:

$$\hat{G}_{M,B}(b, b; n) = \frac{1}{nT_n h_G} \sum_{t=1}^{T_n} \sum_{i=1}^n K \left(\frac{b - b_{it}}{h_G} \right) 1 \{ m_{it} < b, n_t = n \}, \quad (4)$$

where T_n denotes the set of auctions with n bidders. What is the analogous equation for $\hat{g}_{M,B}(b, b; n)$?

6. Suppose that we have two types of bidders, A and B , where members are identical within each class. Now the composition of bidders in an auction matters, not just the total number. Modify the two estimating equations in the first step to accommodate this heterogeneity.