Optimal Dynamic Auctions for Revenue Management

Gustavo Vulcano

Garrett van Ryzin

Costis Maglaras

Graduate School of Business
Columbia University, New York
gjv1@columbia.edu

Abstract

We consider a revenue management problem in which buyers act strategically and bid for units of a fixed capacity over time. We prove that modified versions of the classic first (descending or "Dutch" auction) and second price (ascending or "English" auction) mechanisms maximize the seller's revenue. Moreover, we show explicitly how to conduct these optimal auctions. We then compare the optimal mechanisms to traditional list price mechanisms. In several important cases, we show theoretically that the optimal auction is in fact no better than using list prices. In other cases, numerical results show that significant revenues gains are achieved by using an optimal auction mechanism.(This paper is an abreviated version of a paper by the same title submitted to *Management Science*.)

Revenue management traditionally involves segmenting customers, setting prices and controlling capacity to maximize the revenue generated from a fixed capacity. See McGill and van Ryzin [9] for a review. However, especially with the rise of the Internet, many industries have lately begun experimenting with alternative pricing mechanisms, such as auctions, guaranteed purchase contracts, group purchasing, etc. (See van Ryzin [15].) Given the capabilities that on-line channels provide, such innovations are not surprising. But they do raise some important theoretical and practical questions. In particular, exactly which mechanisms should be used to maximize revenue in any given context? How should an optimal mechanism be designed and operated? And how much benefit (if any) can be obtained from a better mechanism?

We make some initial progress in addressing these questions. Specifically, we consider a natural variation of the traditional single-leg, multi-period revenue management problem (e.g. see Brumelle and McGill [1]) in which buyers bid strategically in each period for a limited supply of capacity. The assumption that buyers act strategically is in contrast to traditional revenue management models, where demand is independent of any decisions made by the seller and buyers do not act strategically. Our model instead closely follows the assumptions of classical auction theory as described in the seminal work of Vickrey [16], the influential paper of Milgrom and Weber [11], the recent survey by Klemperer [5] and earlier survey articles: McAfee and McMillan [8], Milgrom [10], Rothkopf and Harstad [13], Matthews [7] and Wolfstetter [17]. As in this auction literature, we assume buyers have private valuations for a unit of capacity, and they act strategically to maximize their utility (i.e. their value minus the price they pay). As a result, a buyer's behavior is dependent on both the pricing and allocation mechanism selected by the seller and by other buyers' strategies.

To our knowledge, few papers have addressed the link between revenue management and auctions: Cooper and Menich [2] proposed a generalization of the Vickrey auction to sell airline tickets for a network of flights. However, this work does not capture the dynamic decision making feature of our problem. Eso [4] worked on an iterative sealed bid auction for excess seat capacity for an airline, where bidders get instant feedback, including minimum bid suggestions for declined bids. She modeled every iteration as a multi-unit combinatorial auction (see de Vries and Vohra [3] for a survey on this topic).

Motivated by Internet auctions, Segev et.al. [14] deal with a problem similar to ours in which an auctioneer tries to sell multiple units of a product using a multi-period auction. The key difference is that customer bids are exogenous in their model. Specifically, they propose a Markov chain to model bidder behavior. Thus, their analysis does not consider the strategic behavior of buyers and does not employ game theory to endogenize bidder behavior, as we do in our work. Another difference is that Segev et.al. [14] assume the seller precommits to the number of units to award in each period. We do not impose this restriction, and indeed show below that precommiting is suboptimal and that the seller is better off observing the bids first and then deciding how many units to award based on the bid values she receives.

Our model is closer in its assumptions to traditional auction theory. However, it differs from traditional auction theory in that the seller receives bids over time. In particular, we assume there are T periods and in each period t, a new set of buyers bids for the remaining capacity. The seller must determine winners in period t before observing the bids (or even the number of bidders) in future periods. This dynamic feature parallels the traditional revenue management model, in which the seller must determine the capacity to sell in a given period before observing demand in future periods.

For this model, we adapt results on optimal auctions from Maskin and Riley [6] to show that dynamic versions of the first-price and second-price auction mechanisms maximize expected revenues for the seller. However, the optimal mechanisms are somewhat more complex than in the traditional auction setting. In particular, in the first-price case, the seller solicits bids, sorts them, infers the bidders' valuations v from the bids (which we show can indeed be done) and then computes what Myerson [12] calls the bidder's "virtual value", defined by $J(v) = v - 1/\rho(v)$, where v is the bidder's valuation and $\rho(v)$ is the hazard rate of the distribution of bidders' valuations. The seller then accepts a bid if its virtual value exceeds the expected marginal cost of capacity. In the second-price case, if the seller chooses to award k, the winners pay the maximum of the $(k+1)^{th}$ highest bid and a threshold price that depends on k and the optimal value function. Under this mechanism, we show it is a dominant strategy for bidders to bid their values v - again bids are accepted if their virtual value exceeds the expected marginal cost of capacity.

What is unusual from a revenue management perspective about these mechanisms that we develop here - but quite natural to auction theorists - is that it can be optimal for the seller to reject a bid even though its revenue strictly exceeds the expected marginal value of capacity. That is, an optimal mechanism, in some cases, will refuse bids that would be strictly profitable if accepted. The reason is that, just as in setting a reserve price in a classical auction, rejecting profitable bids *ex post* is necessary to induce buyers to submit higher equilibrium bids *ex ante*. The expected gains in revenue from these higher equilibrium bids more than compensate for the losses in revenue incurred in cases where profitable bids are rejected. Exactly which bids should be selected and which rejected to induce the revenue-maximizing bidding behavior is given precisely by our analysis. In addition to showing theoretically that modified first-price and second-price mechanisms are optimal, we provide an efficient method to compute the optimal selling mechanism in each case.

We then compare the optimal mechanism to a stylized version of a traditional revenue management settings, which we call a list price, capacity-controlled mechanism - or LPCC for short. The LPCC mechanism sets a fixed, take-it-or-leave-it price in each period together with a limit on the number of units of capacity that can be sold at the list price. The LPCC is equivalent to the static revenue management model, such as the one analyzed by Brumelle and McGill [1], but with prices as well as capacity controls optimized. For this mechanism, it is a dominant strategy for buyers to attempt to buy if the list price is less than their valuation. As a result, buyers' strategic behavior under LPCC coincides with the assumptions of the traditional revenue management model.

We show theoretically that the LPCC mechanism is in fact optimal when there is at most one buyer in each period - or asymptotically as the number of bidders and units to be sold grows large. Our numerical results show that LPCC revenues decrease relative to the optimal mechanism when the number of bidders per period increases (i.e., there is more "aggregation" of buyers) and when the product is relatively scarce. Moreover, this difference grows larger if the variability in buyers' valuations or in the number of bidders per period increases. These results suggest that some aggregation of buyers is necessary to achieve a benefit from using an auction mechanism for revenue management, and that scarcity and more variation in willingness to pay will help as well. Such conclusions are roughly consistent with pricing mechanisms observed in various industries. Indeed, list pricing is typically used for products that are purchased infrequently over extended periods of time and have somewhat low variance in valuation - or products that are purchased in large volumes (e.g. retail purchases of common consumer goods), whereas auctions are typically employed for items that are sold in low volumes and are sold during a single event to a pool of potential buyers who have, plausibly, a high variability in their valuations (real estate, art, bonds).

References

[1] S.L. Brumelle and J.I. McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41(1):127–137, 1993.

[2] W.L. Cooper and R.P. Menich. Airline ticket auctions: Revenue management and the pivotal mechanism. Submitted to *Management Science*, October 1998.

[3] S. de Vries and R. Vohra. Combinatorial auctions: a survey. Technical report, Zentrum Mathematik (TU München) and Kellogg Graduate School of Management, June 2001.

[4] M. Eso. An iterative online auction for airline seats. Technical report, IBM T.J. Watson Research Center, February 2001.

[5] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999.

[6] E. Maskin and J. Riley. Optimal multi-unit auctions. In P. Klemperer, editor, *The Economic Theory of Auctions*, volume 2, pages 312–336. E. Elgar Pub., 2000.

[7] S. Matthews. A technical primer on auction theory I: Independent private values. Discussion Paper No. 1096, CMS-EMS Northwestern University, May 1995.

[8] R.P. McAfee and J. McMillan. Auctions and bidding. *Journal of Economic Literature*, 25(2):699–738, 1987.

[9] J.I. McGill and G.J. van Ryzin. Revenue management: Research overview and prospects. *Transportation Science*, 33(2):233–256, 1999.

[10] P. Milgrom. Auctions and bidding: a primer. *Journal of Economic Perspectives*, 3(3):3–22, 1989.

[11] P. Milgrom and R. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.

[12] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[13] M.H. Rothkopf and R.M. Harstad. Modeling competitive bidding: A critical essay. *Management Science*, 40(3):364–384, 1994.

[14] A. Segev, C. Beam, and J. Shanthikumar. Optimal design of Internet-based auctions. *Information Technology and Management*, 2(2):121–163, 2001.

[15] G.J. van Ryzin. The brave new world of pricing. In "Survey–Mastering Management", *The Financial Times*, October 16th 2000. Pg. 6.

[16] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

[17] E. Wolfstetter. Auctions: An introduction. *Journal of Economic Surveys*, 10(4):367–420, 1996.