

The Customer Valuations Game as a Basis for Revenue Management

Kalyan Talluri *

December 22, 2008

Abstract

I describe the customer valuations game, a simple intuitive game that can serve as a foundation for teaching revenue management. The game requires little or no preparation, props or software, takes around two hours (and hence can be finished in one session), and illustrates the formation of classical (airline and hotel) revenue management mechanisms such as advanced purchase discounts, booking limits and fixed multiple prices. I normally use the game as a base to introduce RM and to develop RM forecasting and optimization concepts off it. The game is particularly suited for non-technical audiences.

Key words. revenue management, teaching, valuations.

1 Introduction

Revenue management (RM)¹ is a very appealing and exciting topic to teach. It deals with pricing, a mechanism that touches all human existence; it applies to industries such as airlines and hotels, whose mysterious practices baffle and fascinate universally; is based upon the latest and greatest information technology innovations; and has strong economic and scientific roots to support it. So one would think teaching it is a hoot. My experience however, teaching it to a variety of audiences (executives, MBA, PhD) in various formats (full course, half course and short sessions), has been more mixed.

In my experience, the standard introductory material, based on a lecture style presentation, would go very well with executives or managers in RM industries such as hotels, who see the problems every day and immediately catch on to what the scientific part of RM is trying to do, but the same material to a general audience would get a lukewarm reception. They all would have *experienced* RM, at airlines or hotels, but do not see it worth their time to know how actually it is done.

At some point I realized that what is missing is that they do not experience the same thrill and tension that a typical revenue manager does in controlling sales and watching his performance *ex*

*Kalyan Talluri, ICREA and UPF, Universitat Pompeu Fabra, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain, email: kalyan.talluri@upf.edu

¹We assume the reader is familiar with the practice and motivation of RM. For background material, please see the textbooks [2] and [4].

post. This is a pity, because a casino-style thrill on winning and making money cleverly is what makes the practice of RM so exciting and surely if one could convey that to the participants they would be hooked. So this is what led me to develop the customer valuations game and I have found it to be a very effective vehicle to carry almost all the basic concepts of RM into the students' minds.

The principal goal of the game is to make the students experience the day-to-day task of a revenue manager. All other learning objectives are somewhat secondary. Specifically, I do not use the game to proselytize a model of consumer behavior or dynamic pricing or revenue management; rather when later if I describe a model for a technical audience, I refer back to the game to support or criticize the model.

The other pedagogical objectives of the game are:

1. Move students away from cost-based pricing and start them thinking of how customers value the product, and subsequently value-based pricing².
2. Make them understand the difficulty and futility of guessing the valuation of the next customer in an anonymous sales process.
3. Start them thinking of somehow correlating customer valuation with some observable signal, i.e., developing screening devices.
4. Motivate airline RM type inventory control where prices are fixed and the revenue analysts control the quantities sold at different prices.

In what follows, I describe the game, my experience with each stage of the game, pedagogical motivation for each stage, and the typical student reactions and questions to expect. Then, I discuss various supportive modules that one can do around the game, constantly referring back to the insights obtained previously.

The game is quite loose and flexible, and allows the instructor to vary the market conditions to emphasize one aspect or the other of RM. I consider this one of the virtues of the game and ask the reader to concentrate on the *idea* behind each stage of the game more than *my* specific implementation of it. Also, students tire after a few hours of this, so to make the best use of their time and energy, I usually limit myself to doing the three variations that I describe here. Doing more variations is of course possible depending on the time and class profile.

2 Game set-up and customer valuations

The idea behind the game is to simulate a stream of customer arrivals and their purchasing decisions, and then switch over and simulate how a revenue manager controls the sale over time of a fixed amount of inventory. Students first play the role of customers and then switch over to the role of the revenue manager. Using the student's own data (specifically, their valuations for the product) enhances involvement and credibility of what they are doing and starts them thinking of customer purchase behavior, valuations and the firm's pricing process. By playing the role of the customer as well as the firm, the student gets to experience and understand the entire sale process.

²Cost-based pricing or margin-pricing is when price is set to the cost plus some margin; while simple and intuitive, "cost" is not well defined when fixed costs vastly dominate marginal costs and demand is highly variable. Value-based pricing essentially ignores costs and sets prices based on customers' maximum willingness-to-pay. The difficulty here is that willingness-to-pay is private information that customers do not normally reveal.

2.1 Game set-up

Setting: A class of between 10 and 20 students that can be divided up into 4 to 5 groups at a later stage. For a larger number of students and groups the instructor should rely on using spreadsheets.³

Props: A set of 20 to 30 index cards convenient for shuffling. A flip-chart or a blackboard, although I prefer the former to recover prior simulations.

Setup: An appealing product is up for sale. It is good to choose a product that has a wide range of valuations, is rather unique, and difficult to price. I use a cruise package to an exotic island and describe it using photos and video. The stay is for a fixed period of time (say 5 days; this avoids questions about length-of-stay controls), all-inclusive, and students are to imagine that it would be under ideal conditions, when they want, with a partner etc. The cruise yacht is small, with around 30 cabins, and the general ambience is of a very high quality.

Valuations step: Students have to write down on an index card the maximum they are willing to pay for the product. I emphasize that they should do this without consulting their neighbors, and that this is the *maximum* they are willing to pay. The valuation is for the entire package, so in my example, it would be for the entire 5 days.

2.2 Customer valuations

There are usually many questions about the product and valuations. Questions about the product are of course specific to the example I use, and for this cruise-package example: Does it include airfare? (*no - as the participants may originate from different places*); What does the all-inclusive part mean? etc. Many of these are immaterial to what follows but can serve to give a pleasant start to the program.

Questions on how they should arrive at a valuation number are a bit trickier however. Should they take the firm's costs into consideration? (*I say no*)⁴; competition prices? (*yes, if they wish*); what they paid before for a similar product? (*again, it is up to them*); their budgets? (*yes*). What I repeatedly emphasize is that they should be honest about what they write down (that this is just a game!).

Now, I collect the cards and if the class is small mix them up with a few other cards from a previous class; it is good to have at least 20 such cards. The pack of cards serves as the customer population. Note that the whole population size is less than the capacity that I announce, so at this stage, capacity constraints are not an issue. Students are aware of this.

I shuffle the cards between one simulation and the next, and cut the set randomly into two, one part of which will be the set of customers for a specific day, to mimic the random size of the market. The instructor can make a deliberate choice on the percentage of cards to choose. For instance, in a pack of 20 cards, should he cut somewhere close to 4 once in a while or always keep it in the range of 15 to 20? I try to keep the number of customers I pick to be around 70% of the total population, as otherwise the simulations end too soon and do not give too much information about the valuations. On the other hand, occasionally choosing a very small number emphasizes the randomness and unpredictability of demand, and is closer to a realistic situation. So I try to do at

³One of the referees has tried the game with 30 students and 8 groups and reports considerable success in implementing the game.

⁴This is a good time also to discuss costs and conclude that many of the costs are actually fixed costs, and the marginal costs are really rather low; so low in fact we can assume they are zero.

least one simulation run choosing a deliberately low number of cards.

The students should now play the role of the revenue manager. They break up into four or five groups. The set of customers whom I randomly cut from the pack of cards arrive over time, and the revenue manager's job is to control the sale of inventory over the sale period.

In what follows I will assume there are four groups.

3 Static Pricing

The first simulation is a very simple game of setting a price, and observing customers arriving over time and making purchase/no-purchase decisions. The firm tallies revenue in the end.

Setting prices is a near universal aspect of commerce but many products have a fixed posted-price; so by starting off with the relatively simple task of setting a single price, students practice their pricing skills and start getting into the spirit of the game.

3.1 Pedagogical goal

The pedagogical goal of the first set of simulations is to understand the difficulty of pricing, and experience how the revenue manager forms an idea of customer valuations. The key objective of this part is to move the students away from cost-based or margin-based pricing to value-based pricing.

3.2 Description

Each group should set a price for selling the package. The price is for the entire package. The price will remain fixed throughout the sale horizon. Since they know the customer population size is less than the capacity, running out of capacity is not an issue. Their objective is to maximize the total revenue collected during the sale period.

Customers (the index cards) arrive over time and each customer buys if the number on the card is greater than or equal to the price, and if the number on the card is less than the price, the customer does not buy the product (finds an alternative and disappears from our system—that is, we are assuming impatient non-strategic customers). I pull out one card after another and note down sale/no-sale (say as A for Accept, and R for reject) in four columns on the blackboard. So students know if a sale occurred or not, but they do not know the valuations. I emphasize here that in practice few firms are able to observe no-purchases, and mention that there are statistical procedures to correct for this lack of information.

I also emphasize that there is no competition between the groups here. That is, the customers do not choose *one* of the groups at the price the group sets, but rather that these are four separate simulations at different prices with the same set of customers.

After each simulation students tally up the revenue they obtained for their group. Table 3.2 shows a sample simulation run with four groups.

	Group 1	Group 2	Group 3	Group 4
Prices:	500	1200	800	5000
	A	R	A	R
	R	R	R	R
	A	A	A	R
	⋮	⋮	⋮	⋮
Total Revenue:	7500	12000	9600	10000

Table 1: Sample data for a simulation run with static pricing. A is for a sale (accept), R is for a no-purchase (reject).

3.3 Analysis

I run around two simulations to let the students familiarize themselves with the process. Typically I find a group pondering and calculating costs in order to price, and I use it to lead them into a discussion about fixed and marginal costs. Students quickly convince themselves that the fixed costs are actually sunk costs, and the marginal costs are really too low to be of any help in pricing. So I ask them again to ignore this aspect and just to base their pricing on their perception of how customers value the product. The group members reveal their own valuations to each other, so their initial price is usually some function of that (an average or some dominant member’s point of view).

The first round prices usually vary by large amounts but they get more reasonable in the second round. Even though they have not seen the majority of the valuations students pick up the general range as they would have seen eight price points over two simulations. They start seeing the randomness in the valuations and the market size and may even stumble upon a very high price winning in one of the two rounds.

4 Dynamic Pricing Game

Now I tell them that they can change the price anytime, any number of times.⁵

4.1 Pedagogical goal

The learning objective of this part is to start the students thinking about correlating customer valuations with signals, and mechanisms to screen customers by their valuations. Students know that capacity is not an issue, but the fact that they can change prices at any time motivates them to try their luck changing the prices around. They naturally discover the key informational requirement of dynamic pricing—the ability to forecast as a function of time. They also start seeing its limitations: if capacity is not an issue, and they cannot correlate customer arrivals and valuations with any observable signal (such as time), the power of dynamic pricing is rather limited.

⁵One of the referees of this article, who has tried out this game in class, finds that students were adjusting too many times by small amounts, and recommends limiting the number of times they can change their prices.

4.2 Description

I make a show of shuffling the cards and cutting the cards to remind the students about uncertainties in the market. While simulating the arrivals, I keep the pack of cards yet to come hidden, so they do not guess how many customers are remaining.

I run two to three simulations of this. After the first round, I ask them to start thinking of strategies rather than just setting a price.

Also, during the first round, I deliberately do not announce the time till the end of the sale. A number of groups would at various points ask how many cards are remaining. I tell them that no firm really knows this, and they shall not either. Then someone or the other points out (if not, I myself volunteer) that on the other hand, the firm knows how many days of sale are remaining, so I start announcing this in the next round of simulations, based on a rough feel of the number of cards left in my (hidden) hand. After every customer, I let them know how many days are remaining till the cruise departure date. I try for a linear arrival rate, but of course as I am doing this manually it appears quite random.

Tallying revenue is a bit more complicated here, so the instructor should make sure he writes down the price changes of each group. Table 4.2 shows a sample simulation run with four groups.

	Group 1	Group 2	Group 3	Group 4
	400 A	1250 R	800 R	5000 R
	A	1200 A	A	R
	800 R	R	R	R
	A	A	A	4000 A
	⋮	⋮	⋮	⋮
Total Revenue	12450	11890	10225	13000

Table 2: Sample data for a simulation run with dynamic pricing. The numbers in each column are the prices posted by each group for subsequent customer arrivals.

4.3 Analysis

Students get very involved during this part. They try out various strategies relying on their gut feeling, switching back and forth on their pricing; some follow a consistent low-to-high strategy and others high-to-low; but in the end, when they look at their revenue numbers, they find that they are not significantly different from the first set, despite their learning from the first round, and despite their ability to change prices.

I try to ferret out from them why they had such difficulties. Even though they could charge a new price for each customer, the fact that they do not know anything about the next customer, and what the customer's valuation is, and how they arrive over time, makes dynamic pricing a wild gamble that they are as likely to miss as get it right. The instructor can ask the following question at some point: Are they setting a price or are they trying to guess the valuation of the next customer? Students are easily convinced of the futility of the latter. I ask them to think about what they can do about this.

At some point during the game or the subsequent discussion they all realize that even though

they have had a feel for the valuations distribution, they need to (a) estimate market size, or forecast future demand (b) they need to understand how customers arrive over time. Without doing these, they might as well have fixed a constant price.

Indeed, if capacity is not an issue there is no real need to change prices; each customer's valuation is a draw from the same distribution and one has to think of probabilities and expected values over repeated plays of the game rather than guessing the numbers in a specific instance of the game. Two things can motivate us to change the prices (in this setting). One, capacity may become an issue and we would want to raise the price to the optimal capacity-constrained price; knowing the number of remaining days can help us in adjusting our forecast of market size, taking high initial demand as a sign of large market size. Two, customer valuation distributions can change over time, for instance if we believe high-valuation customers tend to arrive closer to departure; again, knowing the number of remaining days help us in deciding if we need to change the price.⁶

Whenever I see a group panicking at the last minute and lowering prices, I pause the game and discuss long-term customer behavior and expectations and the dangers of last-minute offers. Note that I do not criticize the students for lowering prices at the end (or indeed any strategy they follow), as indeed that may be the right thing to do. I just use it to discuss the long-term danger of strategic customers, or customers forming expectations of lower prices near the end and delaying the purchase.

Also by now cost based considerations would seem like a distant memory. All the focus would be on what is written on the cards, how many days are left, and how much capacity we have sold till now—exactly what a typical revenue manager would focus on.

5 Post-screening Game

Right after the dynamic pricing game⁷, I start a discussion on screening mechanisms, RM product design and the general task of correlating customer valuations with some observable signal such as time or purchase characteristics. I discuss other screening devices commonly used in practice: Saturday-night stay, cancelation penalties, age etc.

5.1 Pedagogical goal

The objective of this game is to get a feel for how quantity-based RM is practiced. Rather than fix prices and introduce booking limits, I use the experience from the first two games to show how setting a menu of prices and setting booking limits is also quite natural and mimics their own intuitive price-setting behavior. Students will see the value of forecasting the size of each segment and controlling the number of rooms they are willing to sell to the low-valuation customers, which

⁶A referee has pointed out that one can take this discussion further to explore issues such as the use of data to update beliefs, filtering the data to give better forecasts, reacting to noise rather than the signal etc.

⁷One could also do a variation of the dynamic pricing game where there is a limited capacity, say 10 rooms in the yacht. In theory dynamic pricing would work better, and this would also force the students to vary price as a function of time and remnant capacity. However, I found the capacity-constrained dynamic pricing variant of the game to add relatively little value considering the time it takes to do two or three simulations. Given the limited history and the limited number of simulation runs one can do, it is hard to convince them of the power of dynamic pricing in this setting. Moreover at this point students start to tire, and I would rather use their remaining energy to cover screening issues and quantity-based revenue management common in airline and hotel industries. Of course, it is up to the instructor to vary the combinations and find the right scenarios that work for him or her.

helps a lot in explaining the RM sequence of forecasting demand and optimizing booking limits.

5.2 Description

I tell the students that either naturally or by inducements, high valuation customers can be identified (imperfectly) from low valuation customers. The reasons I mention are that perhaps price-sensitive customers are willing to risk committing their vacation plans well ahead of time for a better price (I mention leisure customers vs. business customers in the airline and hotel contexts); whereas, less price-sensitive customers prefer waiting till their plans firm up, etc.

So the game starts *after* these exogenous screening devices are put in place and low valuation customers have naturally gravitated towards earlier bookings. Advanced purchase discounts are an incentive the firm can offer to make this shift if there is a correlation between valuations and cost of shifting for the consumers, but this game is simpler; here I assume that by consumer expectations of past prices or fear of being rationed, the shift has already occurred.

The challenge here is even after perfect segmentation, we need to price and control inventory, the operational task of pricing and RM departments. I announce that the yacht has only 10 cabins (it has become even more exclusive) and they have to keep this in mind.⁸

Using the same index cards, I split the pack into two based on the valuations, and tell them that one group is made up of low-valuation customers who arrive first. I state that all customers in this group have valuations below a certain number (and reveal this number to make their job easier), and the other group is made up of high-valuation customers with valuations above this cut-off number.

I shuffle and cut both the groups separately, to mimic the random market sizes of each group. The number of cards in each group is a function of the distribution of valuations and the cut-off point that I choose. I try to choose a cut-off point so that there are at least 6 or 7 customers in the high-valuation group (if I am dealing with a customer population of 20). For instance, in my most recent implementation of the game, the valuations ranged from 80 to 10000 euros, and I set 1200 euros as the cut-off point to separate the cards, yielding 7 customers for the high-valuation group.

The task remains the same as in the dynamic pricing part, but students have to keep in mind now that there are fewer rooms, and there are two distinct segments with different valuation ranges. I again make a show of hiding the set of cards to come so they do not see how many customers of each segment are yet to come. To make the game simpler, I announce when the low valuation customer segment is exhausted.

As before, I announce how many sale days are left, more or less as a function of how many cards are left, roughly trying for a linear arrival rate. The whole thing can get a bit tricky physically: pulling out a card, hiding the set, and writing at the same time, so an assistant may come in handy here.

⁸One could also do two variations here: one where there is no capacity constraint, but there are two perfectly segmented groups, and the other with capacity constraint. In the interest of time, I just do the latter. The problem with the no-capacity-constraint scenario is that it becomes, under the perfect segmentation conditions and perfect self-selection that I had described, similar to two independent dynamic pricing games, and consequently does not add much to what has already been done.

5.3 Analysis

I run two to three simulations of this. Students—quite logically—follow a strategy of low prices for the low-valuation segment and high ones for the high-valuation segment, although an occasional group might price only for the high-value customers. The surprising thing is that most groups naturally gravitate towards setting just two prices, one for each segment. They take a few bookings at the low price and then move the price up for the high value segment. One possible explanation is that by now they have seen the same customers over many simulation runs at various price points, and form an idea of the winning prices. As their ability to forecast customer arrivals as a function of time is limited, they stick to a constant price policy based on the distribution of valuations. Of course some groups feel they have not got this price right and adjust the prices as they go along.

The simulation outputs from this last set of runs are very valuable as I use them to introduce two important RM concepts later: the role of forecasting and booking limits. As at various points, during this or the previous stage, some group or the other would have asked to know how many cards are left, I introduce forecasting with the line “wouldn’t you have liked to know how many cards were left?”

As most of the groups end up following a policy of taking a few low-valuation bookings and jumping to a higher price, I use it to introduce booking limits. The fact that in airline type RM, prices are fixed for each segment should seem logical and natural to the students at this point.⁹ Since we cannot guess any individual customer’s valuation within a segment, and since we do not have enough information or history to forecast demand as a function of time, we might as well fix a single price for that segment based on our feel for their valuation distribution.

6 Auxiliary modules

The valuations game can serve as the foundation for developing the course along various lines. I follow it up with three main modules.

6.1 Forecasting Module

As most students during the game wish they had seen how many cards are left to come in the simulation, it is natural to introduce forecasting from there. I tell them rather than being passive about it, one can try to use data, seasonality and knowledge of the market to make a forecast of how many low-valuation and high-valuation customers there are for a specific date in the future.

From this point, the level and background of the class determines what I cover. If it is a non-

⁹Setting a constant price for each segment and controlling the capacity allocated for the lower segment, is roughly inline with the dynamic pricing literature, and is a reasonable heuristic. For a single segment with a fixed capacity, it is well-known that a constant price policy is not too far off from an optimal dynamic pricing policy with stationary distributions (to give just one example, Gallego and van Ryzin [1]). So when the capacity constraint is not an issue and there is perfect segmentation (perfect self-selection to the two distinct time periods) we are solving two separate dynamic pricing problems, and two constant prices would also be close to the optimal. With two segments and a capacity constraint, the heuristic solution has three decision variables, the maximum capacity you are willing to sell for the low customers, and once this is fixed, the constant prices corresponding to the capacity allocated for each segment.

This heuristic can be considered as the two-segment version of the constant price policy. Moreover, we do not need to forecast as a function of time as in dynamic pricing which makes it very recommendable in a class-room setting.

technical class, some simple Excel based exercises on filtering the data and using the average and standard deviation of the data as the forecast. Or, taking it to a slightly higher level I end this module with an introduction to forecast errors, risk and some simple time-series methods. If it is a technical class I do forecasting at the level one sees in an introductory Operations Management class.

6.2 Optimization Module

To introduce quantity-based RM optimization, I refer to the post-screening game. Here, as I mentioned earlier, most tend to accept a few bookings at a low price and then raise the price when they capture a few. Most also can easily be convinced at this point that it is futile to guess the next customer's valuation and that in the absence of reliable data on arrival patterns of customers, fixing a constant "optimal" price for each segment is not such a bad idea. So this serves as a great way to introduce a fixed menu of prices, booking limits and convince them of the need to optimize booking limits—something that otherwise would have seemed like a very odd practice indeed.

For the optimization, again I follow the level of the class. At the most basic level, I try to convey Littlewood's rule and expected marginal values using an Excel spreadsheet. For a more advanced student group, I take it up to network optimization models and algorithms.

6.3 Software Module

Up until now everything has been pen and paper, and perhaps a bit of Excel. I think it is fundamental to convey to the students that a big difference between RM and standard pricing is that it is operational, model-based and designed to be done for the large part automatically by algorithms.

For this purpose I find it absolutely necessary that the students play with some software and see a few screens where the reports and charts are shown connected to a database. I would like to use a student version of some existing software, but I don't believe there is any such thing available. There is a hotel front-office simulation package available (Amazon.com) but the RM part of it does not quite convey what is done in my class.

So far I have used two modules with reasonable success. One is the package developed by Ioana Popescu at INSEAD based on the Easy Profit case [3] and the other one is a Java based module on the internet developed by Larson and Odoni as part of their Urban Operations Research book: http://web.mit.edu/urban_or_book/www/animated-eg/ym/page1.html. The latter, at the time of this writing, is rather cranky, so the instructor should figure it out well before trying it in class (Tips: the tab button has to be hit at just the right speed; there is no other alternative than to close *all* browser windows and restart the program when it freezes, which it will do frequently).

7 Conclusion

I believe that the valuations game is a very good starting point for teaching RM at any level. I find the game particularly useful for students with little or no technical background, or a general MBA class. It can play an important role even in a technical class, serving to instill the fundamentals of the real problem before students get influenced by the incumbent models. It is also a great leveler in classes with mixed backgrounds as, at least at this stage, differences in technical levels tend to be

immaterial to the student's success at the game—so even students with good statistical knowledge appreciate the difficulty and subtlety in pricing.

8 Acknowledgments

Valuable input from Lluís Colomines of Aïda Diodeplus is acknowledged. In addition I would like to thank two anonymous referees and the associate editor for detailed and constructive comments. One referee in particular tried the game in class and suggested many variations of the implementation and ways to improve this write-up, and gave very detailed descriptions of his experience from playing the game—for all of which I am very grateful.

References

- [1] G. Gallego and G. J. van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40:999–1020, 1994.
- [2] R Phillips. *Pricing and Revenue Optimization*. Stanford Business Books, USA, 2005.
- [3] I. Popescu. Easy profit, a revenue management pilot. Technical Report 03/2006-5341, INSEAD, Fontainebleau, France, 2006.
- [4] K. T. Talluri and G. J. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer, New York, NY, 2004.