Dear All,

Thank you for accepting to read a new version of the paper you read in January.

This new version is still very much a work in progress, and is rather incomplete I am afraid. I put a lot of efforts in the framing. I felt this was what I needed to clarify at that stage of my research. The framing section (I.) became quite overdeveloped, at the expense of the story itself, which needs to be completed, and finished (there are missing parts in the end).

My aim is to write an article, one that could be submitted to a history of economics journal. Therefore, I would welcome your comments:

- on the framing: does it work, does it need to be partially or completely changed?

- on the narrative: which parts of the story are unclear, or need to be completed? What is missing, what needs to be developed or removed…

But obviously any comments on the project, and on how to develop it further, will be much appreciated.

Many thanks in advance,

Guillaume.

‘We Want to be Intelligent About Where We Take That Profit’ Science, Technology and Business in the transition from Revenue Management to Marketplace Engineering

Guillaume Yon
guillaume.denis.yon@gmail.com

HOPE Lunch
April 21st, 2023

"In the aggregate, all these mechanisms accomplish roughly the same task: they uncover prices at which buyers and suppliers are matched. It is even plausible that, in the aggregate, over the long run and on average, the prices produced by these mechanisms are, approximately, the ones predicted by classical economic theory. However, such general statements are cold comfort to participants in this turbulent game of matching supply with demand. The trick for them is to deftly surf the market. And that is where technology comes in."
Garrett van Ryzin[1]

I.

In a wonderful piece written for the 65th anniversary of the journal *Management Science*, Ilan Lobel tells the story of the field called 'revenue management'.[2] Lobel is himself a practitioner of the field, from the younger generation of researchers.[3] His audience in this piece is his colleagues in the very large and multifaceted field of management science and operations research. His explicit aim is to make them aware of a broadening of the scope of the discipline of revenue management: Lobel explains how revenue management, from the 1980s to the mid-2000s, was demand management in two main industries, airlines and hospitality. By demand management, Lobel means the use of algorithms to manage the interface between the firm and the market, at that time mostly through controlling quantities sold (seats allocated to different fare classes for instance) or prices, with the aim of maximizing revenue. According to him, somewhere in the mid to late 2000s, the field moved from demand management to market management. Now algorithms manage whole marketplaces, or various aspect of supply and demand, in real time, in new industries like online retailing, online advertising, and ride-sharing.

Laying out neatly the story of revenue management is not the only achievement of the piece, as Lobel also proposes a way to understand the transformation of revenue management in the late 2000s. He points at a shift in the wider economy, more specifically in the way consumer-facing businesses operate, and he terms this shift 'algorithmization'. More and more business decisions are delegated to algorithms. These decisions are automatized and do not rely anymore on humans. This process has been particularly spectacular in the three industries mentioned above: advertising, retailing and urban transportation. Contracts and fixed prices gave way to algorithms setting prices and allocations in real-time.

The third interesting feature of the piece is the way Lobel writes about how the transformation of revenue management as 'a field' is related to this shift in 'the business landscape'. New business models 'opened up a new space', or 'a new research frontier', and revenue management scholars 'moved in'. Yet this movement also 'transformed the field of revenue management in the process' (p. 5389). In the conclusion of the paper, revenue management is 'at the forefront of a transformation of the business landscape', with 'a strong record of integrating theory and practice, by both generating tools for industry and learning from the latest business problems' (p. 5396).

---

[1] van Ryzin, G. 2000. The Brave New World of Pricing. Survey - Mastering Management. *Financial Times* Oct 16.

[2] Lobel, I. 2021. Revenue Management and the Rise of the Algorithmic Economy. *Management Science* 67(9), 5389-5398.

[3] Currently a professor of 'Technology, Operations, and Statistics' at New York University's Stern School of Business. He earned his PhD (MIT, Operations Research) in 2009 – he is from the younger generation. He has a background in electrical engineering, as many practitioners in the field. His personal webpage on the Stern School website provides the following description for his research: "His research focuses on issues related to operations and new technologies, including issues such as pricing, learning and contract design in dynamic and networked markets." Ilan Lobel wrote important papers on pricing for ride-sharing applications. https://www.stern.nyu.edu/faculty/bio/ilan-lobel; last accessed April 14th, 2023.

The purpose of the present paper is to expand and develop the story told by Lobel about revenue management from the vantage point of the history of science and technology, while trying to historically clarify the nature of the relationship between 'the field' and 'the business landscape'. It is the contention of this paper that this relationship can fruitfully be thought as the co-construction of a business practice and a science, or, if we indulge for a moment in catchy formulations, the co-construction of a way of knowing and a way of making money. More accurately, and more than a co-construction, this paper will try to tell the story of the progressive emergence of a common matter of concern.

The next section provides a first definition of this matter of concern. The following section contextualizes the emergence of revenue management as a matter of concern within broader historical changes described in the literature (history of technology, business history, history of operations research and economics), in order to further specify the historical questions we are asking in this paper, and historically situate the significance of revenue management. The next section presents the main chronological articulations of the story. The last section of this first part is a brief presentation of the sources used to tell the story.

Revenue Management as a Matter of Concern

Matters of concern is a concept created by Bruno Latour to broadly understand sciences in society, or the regime of deep entanglement of techno-scientific products with the social fabric. This process, according to Latour, is not captured by the concept of matters of fact, located in a nature untouched by the social realm, and speaking by themselves with the authority of science. Speaking of matters of fact, according to Latour, tends to obscure associated milieu of instruments, practices, institutions and scientists required for the production and continued maintenance of scientific facts. The notion of matter of concern also helps to get rid of deterministic or ballistic accounts, where 'society' shapes 'science' or scientific discovery 'impact' society. Instead, thinking in term of matters of concerns, one can describe the progressive stabilization of scientific facts and the associated environment required for their survival.[4] In a way, the concept of matter of concern is a development of one of the core tenets of actor-network theory: the localization of 'Science' in an interconnected network of practices, which does not mean that all practices are scientific, but that the process of their differentiation can be empirically studied.

The story told in this paper is more directly inspired by the analysis of specific subset of matters of concern. Harold Cook, In a piece on forms of knowledge and the commercial economy in the Low Countries in the early modern period (late sixteen century, early seventeenth century) writes against the common argument that draw a simple causal link between the Scientific Revolution and the economic take-off of the West.[5] He shows instead that the development of long-distance commercial relationships, and the issue of trust raised by those long-distance commercial relationships, made the Dutch merchants value factual description about things, descriptions that can be communicated across time and space, in others words, the modern facts

[4] Latour developed the concept in many venues. The most often cited and compact presentation is Latour, B. 2004. Why Has Critique Run out of Steam? From Matters of Fact to Matters of Concern. *Critical Inquiry* 30(2), 225-248. However, and even though the concept of matter of concern is not explicitly developed there, for our purpose here the most useful reference is Latour, B. 1999. Science's Blood Flow: An Example from Joliot's Scientific Intelligence. In *Pandora's Hope: Essays on the Reality of Science Studies*. Cambridge: Harvard University Press.
[5] Cook, H. 2012. Moving About and Finding Things out: Economies and Sciences in the Period of the Scientific Revolution. *Osiris* 27(1), 101-132.

of the Scientific Revolution (p.132). Conversely, Cook shows that this form of knowledge supported the explosion of the Dutch commercial economy, and even became a chief matter of exchange, as in medicine and technology, furthering the Low Countries material betterment. From a common matter of concern and set of circumstances, Cook argues, a new form of knowledge, associated with the Scientific Revolution, and an economy, the commercial economy, were coproduced.[6]

I am obviously not arguing that the same relationship is at play in the history of revenue management, nor am I arguing that the emergence of revenue management in late twentieth century United States had the same historical significance than the coproduction of the commercial and scientific revolutions in early modern northwestern Europe. Furthermore, Harold Cook's piece cannot be used as a general theory on the relationship between science and the economy, as there is no reason to believe that such a process of coproduction has any generality. Yet Harold Cook carefully suggests the possibility of the existence of some other cases, or the existence of processes of coproduction that have a family resemblance with the one he investigated.[7]

The story, in this paper, is about the progressive emergence of a matter of concern that can be defined as follows. Revenue management addresses a problem in allocation and pricing. Revenue management is, first and foremost, a set of technological systems whose outputs are decisions on which product to sell, to which customer, at what price, where, and when. These decisions are made in order to increase revenue for the firm. 'Increasing revenue' can accommodate a variety of business objectives and business visions. Yet each vision is reflected in a way to make profit, at scale, and revenue management aims at implementing this vision. There are at least three characteristics of revenue management systems that needs to be emphasized to capture their singularity. First, revenue is increased, and profit is made, because of the very tactical, fine-grained, transaction by transaction, real-time approach that the implementation of these systems allows. Second, these systems operate in consumer-facing businesses in large retail markets. These are mundane markets, and we interact with these systems on a daily basis in the western world: they have been implemented by airlines, hotels, car rentals, online or brick-and-mortar retail, online ads which support the free provision of the applications we routinely use online, like search engines, webmails or maps. Third, a technological system here means at least a forecasting component, gathering data on demand and making sense of these data using a model, an algorithm that makes a fast decision based on these data (the decision mentioned above, which can be on the quantity of the product offered, its price, the offer of a specific assortment of products, or all of these), lastly a device to display that decision to customers. This characterization of revenue management might help to understand it a matter of concern: both a business practice and a form of knowledge, and the purpose of this paper is to show how they were historically co-constructed.

[6] "The rise of European commercial power and science at the same time was no accident. Both appeared on the back of new forms of social life that are described as formal and informal institutions and that made information exchange sovereign." *Ibid.*, p.130.

[7] "Such views might be generalizable to other periods and places, but only if the terms *economy* and *science* are pluralized in order to press otherwise essentialist abstractions toward a set of diverse and interactive processes located in time and place, which in turn might invite further historical inquiry. In other words, the relationships between economic activity and scientific activity will change according to the circumstances that produce them both. The manner in which they are interconnected in early modern Europe may not apply to, say, the nineteenth-century United States of America, or twentieth-century East Asia, or other places and times." *Ibid.*, p.102. Emphasis in the original.

To put it differently, revenue management systems link some operational decisions within a firm with economic decisions at the interface of the firm with the market, i.e. pricing-related decisions. These systems therefore make operational decisions, and are often deeply embedded in the operations of the firm, in order to integrate aspects of the internal operations of the firm with profit-making on the market, using automation. Revenue management systems are in a way cognitive systems, on top of, or nested within, other decision-making systems for the operations of the firm, and this cognitive layer emerges because the firm needs to deploy a sophisticated strategy to make profit and implement this strategy at scale. As one of my informants, who happen to be a major figure in the field, told me during an interview, "we want to be intelligent about where we take that profit". This is, I think, a nice way to capture what revenue managements systems are about, and to point at a matter a concern, as defined by the authors mentioned above: making money in these settings is a heavily cognitive activity (or, in a more corporate language, it requires heavy investments in science and technology), and conversely these systems produced a form of highly operative knowledge geared towards money-making. The historical study of this process of co-construction is the topic of this paper; or, how, when and where this common matter of concern came into being.

Technologies, Business Practices, Operations Research and Economics in Revenue Management.

To frame the story, and to explain further what exactly was co-constructed, it is worth situating revenue management on the backdrop of broader transformations described in the literature. This will help me to refine the historical question being asked, and to define further the extend and historical significance of revenue management as a matter of concern.

From Global Distribution Systems to Online Marketplaces

First, revenue management is a technological development.[8] The technological system itself, as sketched above, is a crucial thread of the story. Airlines needed to take very seriously the booking process, the record keeping, how many seats are available in each flight, and the associated identity of the passengers. The process became a challenge with the increase in the number of passengers. Big chalkboards and small cards were not up to the task. This challenge explains why, as early as the 1950s, American, Delta and Pan Am partnered with IBM to create one of the first commercial use of the computer (i.e. outside scientific research or the military), the SABER system (for Semi-Automatic Business Environment Research). SABER only required booking agents to type passengers' information, the computer updated passenger counts on each flight. This massively increased the transaction processing capabilities of airlines, as keeping track of reservations on one single big board hugely slowed down the process. In the 1970s, under the pressure of travel agents, who played an essential role in selling tickets to customers, this internal booking system became a global distribution system. A new version of the SABER system was installed in travel agency and corporate offices (United installed a competing system, Apollo). Travel agents could see in real time the flights and their availabilities, and request bookings. By the early 1980s more than 85% of all airlines tickets

[8] This aspect features prominently in the book by Andrew Boyd, in particular in his chapter 3 (titled 'The Computer Did It'): Boyd, A. 2007. *The Future of Pricing: How Airline Ticket Pricing Has Inspired a Revolution*, New York: Palgrave. Andrew Boyd spent most of his career at one of the leading vendors of revenue management systems, PROS (for Profit and Revenue Optimization Software). His book is made of recollections from his career, and is mostly focussed on airlines.

were sold through these two systems (SABER and Apollo), in travel agencies. This was the first electronic marketplace, or e-commerce environment.[9]

This is an essential element of context for our story, first because this environment generated data, not raw data, but organized data, historical series of point-of-sale data for each flight: numbers of seats sold, price, date of sale, passenger itinerary – or, to put it simply, data on how passengers flew. This is an important element of context for a second reason: a booking system turned into a global distribution system, what SABER and Apollo became in the late 1970s, was not a pricing system. An additional system is needed if one would like to make dynamic pricing decisions when bookings are piling up. This is the first transformation that our story aims at understanding, accounting for both the technical affordances and the technical constraints offered by these global distribution systems when it comes to using them for pricing.

As Latour argued, matters of concern are first and foremost things, their materiality matters. The first thread of our story is therefore technological, on two levels. First, the e-commerce environment changed over the course of our story. Global distribution systems morphed into online marketplaces, the advent of mobile technology and mapping offered a third new environment – yet the fact that these environments are all electronic marketplaces explains why revenue management invaded the latter two. Second, revenue management systems are systems nested within these environments. Their own technical transformations matter too; for instance, when finding way for revenue management systems to leverage new data, or to display prices differently.

The Management of Heterogeneous Demands

The second element of context, and the second thread of the story, relates to business history. One could argue that the characterization of revenue management presented above (in a condensed way, decisions on which product to sell, to which customer, at what price, in order to maximize revenue) is just the definition of what doing business is. The fact that with revenue management these decisions became automated is a first transformative aspect. However, there is something more distinctive for the history of business practices going on in the history of revenue management.

Business historians have convincingly argued that the 1970s saw the demise of the Chandlerian firm – large, vertically integrated firms, which used managerial coordination to achieve high capacity utilization for large economies of scale, producing standardized products at a low cost. Lamoreaux, Raff and Temin show that two things changed in the 1970s.[10] First, rising income per capita made American consumers develop a taste for variety and more personalized goods, goods that reflect their own individual preferences and lifestyles.[11] The standardized goods, produced at a low cost in large quantities, became less attuned to consumers' preferences. Second, falling transportation and communication costs made transactions costs less of a

---

[9] Boyd, A. and Bilegan, I. 2003. Revenue Management and E-commerce. *Management Science*, 49(10), 1363-1386.

[10] Lamoreaux, N., Raff, D. and Temin, P. 2003. Beyond Markets and Hierarchies: Toward a New Synthesis of American Business History. *The American Historical Review* 108(2), 404-433.

[11] Lamoreaux, Raff and Temin provide many nice numbers on this, for instance: "In bedding, for example, the market share of white sheets dropped from about 75 percent in 1960 to about 20 percent by the mid-1970s, while the share of fancy patterned sheets showed precisely the opposite trend." Numbers are also striking when it comes to clothes, for instance men's white shirts. *Ibid.*, p. 428.

problem, and therefore managerial coordination and vertical integration less needed. Lamoreaux, Raff and Temin point out that the outcome of these two changes was the movement towards "corporate shrinking", starting in the 1970s: firms contracting their boundaries, moving away from diversification, outsourcing (and offshoring) production. But they also allude to another effect of these two changes, particularly visible in retail.[12] The greater variety of consumers' tastes, coupled with the long lead time and large quantities required by the manufacturers for department stores chain, made company's buyers more prone to mistakes, with unsold inventories piling up. New firms, like The Gap and The Limited, were based on shorter lead time and offshore production, coupled with technologies leveraging data on what was selling (style, color, size), in every time period, in every location, this in order to better know demand and adjust orders and assortment accordingly, but also to manage prices, markdowns, and therefore inventories much more flexibly.

It not a coincidence that one of the building blocks of the field of revenue management, Guillermo Gallego and Garrett van Ryzin dynamic pricing paper, published in 1994, was the output of a collaboration with a large department store chain, and that, in their introduction, the two authors discuss in depth the examples of The Gap and The Limited in order to set up their price optimization problem.[13] Kalyan Talluri and Garrett van Ryzin, in an important book for the history of the field, make this clear.[14] "Customer heterogeneity", "demand variability and uncertainty", "production inflexibility" (the existence of inventory, or production lead time) and "data and information system infrastructure" are the four "business conditions conductive to revenue management". They explain that the variety of products sold and the diversity of consumers' preferences and purchase behavior, associated with the availability of point-of-sale data, requires firms, if they want to survive, to "manage demand": to make real-time, fine-grained (i.e. transaction by transaction) decision on which product to sell, to which customer, at what price. If not "significant opportunities for incremental gains – on particular products, at particular locations, at specific point in time – are simply lost". For the two authors "a large retail chain, which can have tens of thousands of SKUs (stock-keeping unit, the lowest level at which we identify inventory – such as men's Arrow blue Oxford shirts, long sleeves, size medium) sold in hundreds of stores with prices monitored and updates on a daily basis", or "a modern large airline, [which] can have thousands of origin-destination pairs, each of which is sold at dozen of prices – and this entire problem is replicated for hundreds of days in the future!" are facing exactly the same revenue management problem. Airlines too were dealing with new types of customers, more heterogeneous, more eager to fly to new destinations, than the standard business customer. According to the two authors, this required a new approach to profit-making.

These remarks shows that customer heterogeneity and lower transaction costs did not only had an effect backwards, with previously vertically integrated firms outsourcing, but also forward. Facing heterogeneous demands, to make new profits, and often to survive, firms needed to engage in managing profit transaction by transaction (depending on the product, the customer, the time and the place). The question is how this forward transformation in business practices

[12] *Ibid*, p. 428. The argument here is based on work done by Raff and Temin on Sears Roebuck: Raff, D. and Temin, P. 1999. Sears, Roebuck in the Twentieth Century: Competition, Complementarities, and the Problem of Wasting Assets. In Lamoreaux, N. ed. *Learning by Doing in Markets, Firms, and Countries*. Chicago: University of Chicago Press, 219-252.

[13] The paper is discussed below: Gallego, G. & van Ryzin, G. 1994. Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons, *Management Science*, 40(8), 999-1020.

[14] Talluri, K. and van Ryzin. 2004. *The Theory and Practice of Revenue Management*. Boston: Kluwer Academic Publishers, p. 5 and 13.

unfolded, but also how it interacted with the first, technological, thread of our story. Doing this requires introducing a third thread, which played a crucial role in articulating new technologies to new business problems, and this thread relates to the history of operations research and economics.

"Solving Business Problems with Computers" and the Tools of Economic Analysis

The story of revenue management cannot be understood without having in mind the "revolution" in American business schools that took place between 1945 and 1970.[15] Four components of this revolution matter for the history of revenue management. First, Augier and March show that after the Second World War, American business schools rooted their activities firmly in quantitative, analytic and model-based methods, away from rule-of-thumb approaches. What was done in business schools moved closer to what was done in the other departments. Second, the science practiced in business schools became interdisciplinary, blending mathematics, statistics, computer science and economics – the exact same blend that constitutes the field of operations research, which found a home in business schools. Third, Augier and March show that the research carried out in business school, using this blend of disciplines, was driven by real-world problems, with the goal of being useful to management. The two authors add that this did not mean directly solving problems in specific contexts. It meant producing research for understanding the nature of the problem, how the problem relates to a more general class of problem, one for which new methodologies for problem-solving have been or can be developed. Fourth, Augier and March show that business schools became not just repositories for best management practices, recording, organizing and disseminating those best practices. Business school became active producers of research and knowledge that aimed at improving business practices. The RAND Corporation – some of its members went on to found or transform business schools – and the Graduate School of Industrial Administration at Carnegie played a pivotal role in shaping this culture of interdisciplinary research on business problems.

Operations research originated in the Second World War, when applied mathematics (optimization), statistics, and economics, were used together to address military problems in a problem setup which required to pay great attention to the computability of the optimal solutions, hence the link with the very first computers and algorithms.[16] The development of some of the core tools of operations research during the Second World War and in the 1950s, linear and dynamic programming, or solution to the problem of inventory control for instance, came from people considering themselves as economists. In the 1970s these tools were used and developed, and these problems tackled, by people in business schools, people publishing

---

[15] "Revolution" is the word used by Augier and March in Augier, M. and March, J. 2011. *The Roots, Rituals, and Rhetorics of Change: North American Business Schools after the Second World War.* Stanford: Stanford University Press. I draw on their findings here.

[16] Histories of operations research tend to focus more on the effect of operations research on economics (Mirowski, P. 2002. *Machine Dreams: Economics Becomes a Cyborg Science.* Cambridge: Cambridge University Press; and Klein, J. 2007. Rules of Action for War and Recursive Optimization: Massé's "Jeu des Réservoirs" and Arrow, Harris, and Marschak's "Optimal Inventory Policy". *Unpublished Manuscript*; Klein, J. 2007. Cold War, Dynamic Programming, and the Science of Economizing: Bellman Strikes Gold in Policy Space. *Unpublished Manuscript*), or on the most abstract aspects of operations research, when, in the 1950s and 1960s, some operations researchers tried to build a general science of rational decision-making (Thomas, W. 2015. *Rational Action: The Sciences of Policy in Britain and America, 1940-1960*. Cambridge, MA: MIT Press; Lemov, R., Sturm, T., Gordin, M. D., Klein, J. and Daston, L. 2019. *How Reason Almost Lost its Mind: The Strange Career of Cold War Rationality*. Chicago: University of Chicago Press). To the best of my knowledge, the work of operations researchers in business schools and in businesses after the Second World War, in the 1950s and 1960s, is less documented, beyond the elements provided by Augier and March.

in operations research and management science journals, all disciplines considered as being outside the realm of economics at that time. What changed between the 1950s and the 1970s was that the sections of economics using mathematical tools became more abstract and theoretical in their aims, or more remote from real-world problem. Operations researchers caring about real-world problems did not relate well to the fundamental projects in general equilibrium theory or game theory. At the same time, the operational needs of the large, vertically integrated American corporation of the post-war era, the creation of well-staffed operations research teams starting to use computers in many industries, including the airlines (for flight and maintenance scheduling, for instance), this vast managerial class in need of training, all this fostered a happy marriage between the operations researchers feeling uncomfortable with moves towards abstraction in economics and business schools. And operations researchers, in turn, participated in the broader movement of transformation of the American business schools, as described by Augier and March.

This is obviously a long and complicated story which I overly simplified, but for my purpose here I only need to establish the fact that, as one of my informants puts it: “In the 1970s, operations research was people solving business problems with computers.” Arguably, this statement does not reflect the whole field, and some quarters of operations research were more versed into applied mathematics, with a loose connection to business problems. However, the point here is that in the 1970s, there was a large contingent of operations researchers in business schools and in many industries, collaborating around the use of computers to optimize the operations of the firm, and these people were considered outside the realm of what counted as ‘proper’ economics at that time. As sketched above, the problem of revenue management, as it emerged in the late 1970s, was a business problem, with a heavy computing component (the first two threads of our story). It is therefore not surprising that operations researchers working in business schools became an important part of the story of revenue management.

Yet the scientific aspect of the story is not only about operations research. As I will try to show, with revenue management, operations researchers, trained to improve the internal operations of the firm (supply chain management, minimizing costs for a given, exogeneous demand) were going to the market, with the explicit aim of managing the market, as they managed the firm – and this helps to understand why they became so successful when, in the online world, firms that were run as marketplaces started to emerge. When moving into the revenue management space, academic operations researchers felt the need to use more and more tools borrowed from economics, in order to compute fully optimal dynamic prices, to bring into the airlines setup more realistic models of consumers’ behavior, or to deal with supply and demand estimation in the ride-sharing context.

The third thread relates to this use of tools borrowed from economics. A growing academic community of operations researchers became involved in revenue management from the 1990s onwards, and had a lasting effect on business practices, while participating in the creation of a field where tools borrowed from economics became more and more central. In this form of knowledge production, the use of tools of analysis borrowed from economics is I believe quite distinctive. I would like to briefly signal what I think are two specificities.

A comparison with mechanism design, a field which is today close to revenue management, might help to capture the first specificity in revenue management’s use of economics tools. Operations researchers engaged in revenue management did not build a market from scratch for policy purposes, to make some aspect of the reality look more like the ideal optimal allocation described in economic theory.[17] They did not start by spotting a market failure caused by some issues in releasing information, in order to then build a mechanism for releasing all the information – usually, an auction – that would fix this market failure. They provided advice to the pricer for the maximization of profit. The systems they improved operates *within* a market – it is not *a whole* market - and make use of many market signals (information coming from the actions of competitors). In the practice of revenue management, the business problem, framed as improving revenue within the realm of actual business and technological constraints, within an actual market, comes first; not a market failure detected thanks to the use of economic theory as an indictment to reality.

Second, in revenue management tools borrowed from economics are not used to carry out an analysis, which will then produce a policy recommendation.[18] The tools are embedded in the technological system, in the system, so to speak. With the intervention of academic operations researchers using tools from economics since the 1990s, these tools became embedded in the routine operations of technological systems for pricing. These pricing systems gave flesh to tools of economic analysis, realized them, by attaching numbers to them in real-time, and making decisions on allocation and pricing based on the actualization of these tools in the system. As one of my informants put it, revenue management systems are "powered by economics […]. It is 'Economics Inside', as in 'Intel Inside'". One could say that revenue management makes economics look more and more like a technoscience.[19] I will try to provides examples of this in the story told below.

Outline of the Story

Thinking in terms of matter of concern helps to frame the historical question in a way that does justice to all its dimensions. The question is to understand how a set of technological systems, a new business practice, and a form of knowledge were co-constructed and shaped each other. There are three moments in this process.

First, from the late 1970s to the early 1990s a business problem emerged in the airlines. The airlines developed global distribution systems prior to the late 1970s, and when deregulation hit in the United States, American Airlines created a new system, called DINAMO, nested into its global distribution system (SABER), in order to manage the quantities of seats associated with different fares classes in each flight. In this section, I present the problem faced by large airlines following deregulation (the threat by new entrants). Then I present the quantity allocation system invented at American Airlines and its optimization methodology. Lastly, I show that, in the late 1980s, the importance gained by these systems in the airline industry, and

[17] As for instance in the well-documented FCC auctions case: Guala, F. 2001. Building Economic Machines: The FCC Auctions. *Studies in History and Philosophy of Science Part A*, 32(3), 453-477. Mirowski, P. and Nik-Khah, E. 2017. *The Knowledge we Have Lost in Information: the History of Information in Modern Economics*. Oxford: Oxford University Press.

[18] As for instance in the use of cost-benefit analysis: Berman, E. 2022. *Thinking like an Economist*. Princeton: Princeton University Press.

[19] This might have something to do with what Mary Morgan describes as the shift in economics to the problem-solving mode of engineering, a shift happening according to her since the 1970s, on the backdrop of the longer story of economics becoming more and more tool-based. Mary Morgan uses the notion of problem-solving mode of engineering to describe a shift in economics: tackling specific and well-defined issues, using the tools of economic analysis, without regards for the derivation from some economic laws. The problems solved remained economics problems. The story of revenue management might suggest that this new use of economics tools for problem-solving travelled further down than the applied within economics, to impregnate practical technological practices in businesses. Morgan, M. 2020. Technocratic Economics: An Afterword. *History of Political Economy* 52(S1), 294-304.

their fancy new optimization approach, turned them into a scientific object. Pete Belobaba from MIT proposed two improved heuristics for airlines revenue management, known as the EMSR (expected marginal seat revenue) decision models. Yet, in this first phase, revenue management is first and foremost an emerging business practice in the newly deregulated airlines, associated with a technological innovation within global distribution systems.

The second phase of the story starts in 1993, and 2005 is a convenient ending point. This is the phase when academic scholars in operations research moved in. In 1993, Guillermo Gallego and Garrett van Ryzin published a dynamic pricing model, a solution with full price flexibility (the computation of the whole price path), when there is a finite inventory to be sold over a finite time horizon. The model was a "full-on optimization", as one of my informants described it, framed as a dynamic programming problem. The paper supported the involvement of the academic community, as it reframed airlines revenue management as part of a generic class of problems, also raised in retail or hospitality, actual business problems that can be tackled with dynamic programming, to create a lot of incremental value for firms. In the 1990s, the model actually helped to solve the so-called network problem in airlines and hotels. When the product sold is constituted of multiple resources (multiple flight legs for an origin-destination flight with a correspondence, multiple nights in a stay in a hotel), the nature of the revenue optimization problem changes. The approach fostered by Gallego and van Ryzin provided a framework to think about that problem. Network revenue management was reframed as a linear programming problem, and the shadow prices were used as controls in actual, implemented at scale, revenue management systems. The 1990s became an era of tight collaboration of operations researchers with two industries (airlines and hospitality), and this collaboration became widely acknowledged and institutionalized. Lastly, I show that around the turn of the century, the growing involvement of operations researchers from academia had another effect. More tools borrowed from economics were introduced into revenue management, new methodologies were developed, in particular for modelling demand. These new methodologies for modelling demand went hand in hand with the new availability of online shopping data. Therefore, around 2005, there were many calls from academics for a major overhaul of revenue management systems, away from the fare class model rooted in legacy airlines booking systems, to embrace full price flexibility and more sophisticated models of customers' choice in forecasting. From the mid-1990s to the mid-2000s, academic operations research moved in the space opened by airlines in the 1980s, transformed business practices, and even went ahead of these.

Yet, in the 2000s, airlines and hotels saw the proposed overhaul as too bold. Adoption of the new approaches for demand modelling lagged. At the same time, new businesses were starting to emerge: online marketplaces. Revenue management entered this new space, and changed in the process. It moved from demand management to marketplace management, the management of the two sides of the market, for revenue maximization. Given the current limitations of my research I am not able to describe how revenue management developed in online retail and online advertising, two important sites of the shift towards marketplace management. I will rely here on the third most important site, the industry on which I have more information: ride-sharing. The third phase of the story sees revenue management morphing into marketplace management, installed at the core of how online marketplaces function.

Data Sources

The research reported here is based on interviews with key figures in the history of revenue management. Interviews provide a way to track the entanglements of science, technology and business practices, as these are only very partially recorded in the published research papers, focussed as they are on the essence of the contribution.[20] Interview also allowed me to identify the relevant primary sources, the key contributions to the field, and to navigate this huge literature.

The developments below rely on 16 extensive interviews of at least one hour, usually one hour and 20 minutes, conducted over Zoom, except for one interview conducted in person, from September to January 2022. One additional interview was made by sending written questions and receiving a written reply. Two key figures in the field have been interviewed twice. Among the 15 interviewees, 2 worked only in industry. The 13 interviewees holding an academic position worked for the following institutions (some had worked for different universities over the course of their career): five at Cornell, three at Columbia, one at Berkeley, one at Stanford, one at CalTech, one at NYU, one at Duke, one at Imperial College London. Most of them worked in the business schools of these universities, but the 3 positions at Cornell and the position at Berkeley were in the engineering school. All interviewees had worked in industry, usually multiple industries, over the course of their career. The academics all did consulting, but they had also all worked as employees, not just consultants. Working in industry could happen before or after an academic career. In most cases, it happened either during a sabbatical, which in the field can span over two to three years, or because they had a dual appointment (in a university and a firm, as for instance with the Amazon scholar status). Three people interviewed worked for firms selling pricing software mostly to transportation companies, three worked for airlines, three for ride-sharing firms, three in online retail, two in hotels, one in online advertising.[21]

I also attended the annual INFORMS conference, which gathered around 6000 attendees in Indianapolis, between the 16th and the 19th of October 2022. Around 180 papers were presented within the Revenue Management and Pricing track. INFORMS is an academic conference, but industry was largely represented, reflecting the overlap between business and science which is at the heart of revenue management. At INFORMS I conducted many informal interviews and had informal conversations with people working in industry and academia, often in both at the same time.

[20] Another source, in addition to the interviews I conducted, and the printed primary sources (research books and papers in revenue management which made the field), is the secondary literature, reviewed in the section above. As the section above shows, there is no historical or social scientific treatment of revenue management, to the best of my knowledge. Yet there are published recollections from actors of this history, and in some research books, practitioners' versions of the history of the field, discussed above.

[21] This list only considers the most significant appointments in industry, just to give an idea to the reader of the industries my interviewees worked for.

II.

New Tactics for Deregulated Airlines (1978-1989)

Framing a New Business Problem: American Airlines vs. People Express

The story of revenue management started at American Airlines after the deregulation of the sector in 1978. Before 1978, airlines were regulated almost as a public service. Any opening or closing of a line, and any price change, required filing a request to the Civil Aeronautics Board (the CAB, abolished in 1985). Prices were controlled, based on some simple mechanism allowing for a guaranteed rate of return for the airline. Likewise, most companies in Europe were state-owned and their prices regulated in the same way. Deregulation was, therefore, a necessary precondition for the airlines to be allowed to change prices quickly, and more broadly to set their prices as they wished, without providing a justification and disclosing their costs to the regulator. However, this is only one aspect of the story. First, factors other than deregulation played an important role in the invention of revenue management in the airlines; second, deregulation shaped the pricing revolution in a more indirect way than just allowing for price changes.

First, as explained above, long before deregulation, American Airlines had partnered with IBM to develop a computerized system for keeping track of flight bookings. It eventually became, in the 1970s, a global distribution system, and the first electronic marketplace: SABER (for Semi-Automatic Business Research Environment, at some point it became rebranded at SABRE). In the 1960s, the SABRE system processed thousands of reservations per hour (customers called operators), making booking information available in real time in a centralized system. In the mid-1970s, SABRE computers were installed in travel agencies, making it possible for travel agents to 'see' the flights directly. The existence of that technology was crucial when it came to manage capacity.

Second, deregulation mattered not just because it allowed 'price freedom' but because, almost immediately afterwards, low-cost carriers entered the market.[22] A company like People Express operated at a much lower cost than American Airlines. People Express used a single model of plane to lower the cost of maintenance, operated only on the most popular routes between big cities, cross trained its staff so that gate agents could serve as flight attendants and had a non-unionized workforce. American Airlines could not possibly start a price war and match the low prices of People Express, nor could American carry on with much higher prices than competition.

Bob Crandall, American Airlines CFO at the time,[23] is credited with the 'surplus seats' idea. Crandall acknowledged that American Airlines was operating at zero marginal cost on most of its flights. Indeed, most of American flights departed with empty seats. The cost of an additional passenger in that case is almost zero. Additional fuel costs are negligible, the number of crew members is fixed and independent of the number of passengers as it is set by safety regulations, and the cost of a drink is negligible in comparison with total costs. Under the threat

[22] Versions of this spectacular part of the story can be found in: Cross, R. 1997. *Revenue Management: Hard-Core Tactics for Market Domination*, New York: Broadway Books; and in Boyd, A. 2007. *The Future of Pricing: How Airline Ticket Pricing Has Inspired a Revolution*, New York: Palgrave.

[23] He went to become American Airlines CEO between 1985 and 1998. He is widely considered as "an industry legend": see the popular treatment of his career in Reed, D. 1993. *The American Eagle: The Ascent of Bob Crandall and American Airlines*. New York: St. Martin's Press.

of low-cost carriers like People Express, Crandall considered these seats as ‘surplus seats’ for potential ‘surplus passengers’. These seats could be treated similarly to those in a flight operated by People Express’, and therefore priced as such. The problem became the identification of the ‘surplus seats’.

According to my informants and the recollections quoted above, Crandall went further and imagined a neatly segmented market. The key invention here is the fencing.

Crandall’s analysis was that People Express was tapping in a market that did not exist in the regulated regime, the leisure travelers: they book early, are willing to stay over for the weekend, are price sensitive. These are families on holiday, couples getting away for the weekend, college students visiting home. American Airlines used to deal with business travelers: they book later, want to be back on Friday night with their families, have little flexibility, i.e. they value schedule convenience the most, and are much less price sensitive because they are ‘on expense’, i.e. travelling with company money.

Crandall designed a way to exploit the existence of these heterogeneous demands: the fences. American created a new fare (the super-saver fare), with a lower price, but available only for those who book early (30 days in advance) and are willing to stay over for the weekend. These restrictions were entirely artificial, they had nothing to do with the cost of the ticket. The aim was the self-selection of the consumers to avoid cross-cannibalization between the two segments (business customers could not abide to the conditions associated with the low fare), thus creating fences.

To put it differently, prior to deregulation, airlines were pretty happy to sell expensive tickets to business customers. With deregulation, new entrants could tap into a market previously ignored. What the fences exploit was only one dimension of the newly acknowledged heterogeneity of demand: time. Leisure travelers tended to book earlier than business travelers, who knew that they had to travel usually closer to the departure date. This was obviously an approximation. Yet it was an approximation which cleverly exploited the technological affordances offered by the SABRE system. In the SABRE system, one could see bookings piling up in time. This approximation created two independent streams of demand within the SABRE system, streams of demand that could be forecasted for each scheduled flight (how many business customers are likely to show up for this flight close to departure is something that could be seen easily in the data provided by the SABER system).

This was a breakthrough because the problem of dealing with the success of People Express, which in the early 1980s was threatening to push American Airlines out of business, was now narrowed down to an optimization problem: how many business seats need to be protected in the SABRE system for each scheduled flight?

A System within a System: DINAMO

Barry Smith, head of operations research at American, and his team, are credited with the implementation.[24] The problem above was a capacity allocation problem, or more technically in the language of operations research, an inventory control problem. The problem is: how many seats to allocate to each fare class? The model developed was static, the allocation was

[24] Barry Smith received with his team the 1991 INFORMS Edelman Prize, for their work on revenue management at American Airlines. See Smith, B., Leimkuhler J. and Darrow, R. 1992. Yield Management at American Airlines. Franz Edelman Award Paper. *Interfaces* 22 (1), 8-31.

decided in advance, for a single scheduled flight. The rule (the algorithm) was as follows: the number of seats in the cabin to be allocated to the low fare was determined by the marginal value of the inventory. One more seat can be added to the low fare class, as long as the revenue of selling one more seat at the low fare was higher than the expected revenue of selling this seat at the higher fare.

This was an operational solution, in two aspects.

First, calculation of the expected revenue required knowledge of the distribution of demand (probabilities attached to sales to business customers). This was easily forecasted in the SABRE system thanks to the self-selection of the passengers created by segmentation. American took advantage of its large number of flights and its repeated sales on each route. Forecasts could be rerun before each new opening of a flight for booking. This was an indirect way to account for changes in demand and in consumer preferences, as these would be reflected in changes in forecasts. Then, inventory control algorithms rerun the allocation of the seats. American could even focus the first forecasts only on something they already knew well: the demand from business travelers. They could run the first allocation algorithms only knowing the probability of the event that the demand from business travelers exceeds the remaining capacity for each possible allocation of seats to the leisure fare. The forecasting component of the system used existing, already available computational resources. This became an essential and enduring feature of the field of revenue management: researchers would pay great attention not just to computability, but to implementation at scale and using the existing technological affordances.

Second, the rule presented above, the algorithm, could be lodged within the existing SABRE system. The control variable in the optimization problem laid out by Barry Smith and his team, i.e. the decision to optimize, was when to close the low fare. Travel agents could log in the request of the customer on their SABRE terminal. The SABRE system just signaled whether one more request for the low fare could be accepted or rejected by the operator based on the seats allocation decided in advance by the algorithm, and the bookings already accepted. This constituted an electronic marketplace, with a – for now very limited – form of price flexibility.

The forecasting component, the algorithm, and the display (displaying the decision to accept or reject the request) constituted a system nested within SABRE, and called DINAMO (for Dynamic Inventory Allocation and Maintenance Optimizer). Two features of this first revenue management system are worth emphasizing, in order to understand subsequent developments in the field.

First, the first American Airlines revenue management system was static optimization. Decisions on the allocation of seats were not revised in real time as demand was piling up for a scheduled flight. The allocation was decided in advance, before opening the flight for booking.

Second, the system featured little price flexibility, only two fare classes, low cost or leisure class and high fare or business class. The price of the ticket for each class was set in advance. This was not because of a technology constraint in the booking system, but because airlines did not want to start a price war. Therefore, they all had roughly the same prices for the same route. Also, airlines published prices in printed press and billboards for advertisement, they needed some price stability. Competition did not mean direct price competition. Competition took place at the tactical level, in the revenue enhancing properties of an optimal micro allocation of the seats between the given fares. And competition was intense. Success or failure was determined by the sophistication and performance of the algorithms. Operations researcher teams would become crucial assets in airlines. In 1985, when the new super-saver fare was launched, American Airlines, then on the verge of bankruptcy, could suddenly capture leisure travelers, while retaining business customers.

Revenue Management became a Scientific Object: the 1987 EMSR Heuristics

The implementation of this revenue management system by American Airlines in 1985 was a tremendous financial success for the company, and it also put People Express out of business in 1987. All airlines in a post-deregulation environment needed a revenue management system. This surely participated in attracting the interest of the operations research community at large. But, at the same time, the business problem opened up by American Airlines was also interesting from a pure optimization point of view. The joint presence of these two aspects, large value creation and contribution to a new approach to problem-solving in general, was the definition of what is a good problem in the field of operations research at the time. An important contribution of the late 1980s, the EMSR model, reinforced the impression that the problem was interesting from a pure optimization point of view, and was a landmark in the structuration of revenue management as an autonomous – and thriving – subfield of operations research.

EMSR stands for Expected Marginal Seat Revenue. The model was developed by Peter Belobaba from MIT.[25] Belobaba was working for Delta Airlines and participated in the implementation of their revenue management system around 1987. His model improved the American Airlines model in two respects. First, it used a more sophisticated probabilistic approach, accounting for uncertainty in demand forecasting. Second, and most importantly, it introduced the idea of nested booking systems. The American model created two completely distinct inventories for the low fare class and the high fare class. Belobaba's nested booking idea retains the booking limit on low fares, but became 'transparent from above' (i.e. from the high fare class), taking advantage of the possibility of accepting high fare class requests when the low fare classes are still open. In the previous model the high fare class was opened after closing the low fare class, creating a strong restriction on the time of the booking. Belobaba's model was more dynamic and allowed revisions of the allocation while the cabin was filing up. The model allowed to move from the two-fare class model presented above to a n-fare class model, and as a result in the late 1980s, airlines could optimize for bookings made up to 300 days in advance, opening and closing up to 26 classes.

The Scientists Move In: Dynamic Pricing, Network Revenue Management, and Choice Models (1994-2005)

The airlines episode attracted a lot of attention, and the EMSR heuristics attracted even more attention within academic circles. The 1990s were a decade when operations researcher took seriously revenue management as it emerged in the airlines in the 1980s. They subsequently made important contributions which improved business practices, in particular around the solution of the so-called network problem.

[25] It was the topic of his (very thick) PhD dissertation: Belobaba, P. 1987. *Air Travel Demand and Airline Seat Inventory Management*. PhD Dissertation: Massachusetts Institute of Technology. The foundational paper published out of that dissertation is Belobaba, P. 1989. Application of a Probabilistic Decision Model to Airline Seat Inventory Control. *Operations Research* 37(2), 183-197.

This was, in a way, a renaissance for operations research: operations researchers who worked on revenue management systems in industry and academia mingled,[26] and were seen as incredibly successful within the profession, developing new optimization techniques that spectacularly improved the bottom lines of many companies. The CEO of these companies not only agreed in giving credit to operations research for these improvements, they found themselves under pressure from financial analysts, who, in the early 1990s, opened annual shareholders meetings at Delta or Hilton with offensive questions about the companies' strategies to overhaul their revenue management systems. Poaching the best researchers and their students became current practice. For operations research considered as an academic discipline which, over the previous two decades, had seen its core client, the military, slowly withdrawing its support, followed by the National Science Foundation, this was very good news. Moreover, the importance of the revenue management systems in these industries created a new market in business schools for the operations researchers working on those systems. In the early 1990s 'revenue management' became a well-established and well-regarded subfield of operations research.

Yet this influx of academic researchers also had another effect, in the same decade, in addition to widely acknowledge improvements to existing systems. More speculative developments took place within the field of revenue management, around dynamic pricing and consumer's choice models. By more speculative, I mean that both with dynamic pricing and models of consumers' behavior, deeper transformations to the airlines and hotels booking systems were required. Around 2005 calls for "ditching old airlines legacy systems", as one of my informants put it, multiplied in the academic quarters of the field. Two dimensions are interesting here.

The first is that one of the characteristics of the academic contributions of the 1990s was to make use of tools borrowed from economics: grounding the network problem is shadow prices analyzed as opportunity costs (therefore, what a revenue management system should do was to constantly actualize opportunity costs based on forecasts and bookings piling up), and improving forecasts with models of consumer's demand imported from developments in econometrics. This "more economicsy outlook" or "full-on optimization", as two different informants put it, were not present in the first system developed at American, and can be considered as a product of the involvement of academic researchers.

Second, these academic contributions went hand in hand with the new idea, at the turn of the century, of leveraging online shopping data. This also was alien to the airlines revenue management systems, and required a pretty big technological change. Yet this matters when it comes to understand why revenue management invaded online marketplaces in the third part of our story: many quarters of revenue management on the academic side were ready for it.

The 1993 Gallego and van Ryzin Dynamic Pricing Model

A first crucial academic contribution was the 1994 paper by Guillermo Gallego and Garrett van Ryzin.[27] The paper was motivated by research done in retail for Federated Department

---

[26] The Airline Group of the International Federation of Operational Research Societies (AGIFORS) was – and still is – an important platform where the entanglement of industry and academic research in revenue management was institutionalized.

[27] Gallego, G. & van Ryzin, G. 1994. Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons, *Management Science*, 40(8), 999-1020. Guillermo Gallego earned in 1988 a PhD from Cornell and was an assistant professor in the department of industrial engineering and operations research at Columbia. His PhD was in operations research, with a minor concentration in economics. He had a dual training in mathematics and

Stores (a large conglomerate that included Macy's), from which the two authors secured a grant. The research was on markdown optimization, a topic close to inventory control, in which the authors were trained. The authors also had first-hand knowledge of airlines' revenue management systems. Their 1994 paper developed a new approach to retail pricing, and also a new framework that could help to consider revenue management (in airlines and hotel) and retail pricing together, as part of a larger class of optimization problems. They christened this larger class of problem: dynamic pricing.

One of the contributions of the paper, and of the large literature spurred by it, is to show that revenue management systems were unnecessarily idiosyncratic. First, because they did not make full use of price flexibility. If the number of fares was not restricted to two in the early 1990s, as in the original American model of the early 1980s, it remained limited. Second, airlines and hotels revenue management systems were rationing systems, based on an allocation mechanism, using the fare classes to ration demand with the aim of maximizing revenue. Using prices to shape demand and adapt it to a fixed supply seemed intuitively more profitable.

The retail problem the two authors had to solve had much in common with the airlines setup. In their retail setting, garments required eight months to be assembled overseas, but the sales horizon was nine weeks. This short sale season was the new rule in the fashion industry in the 1990s; the authors mention that The GAP worked on the same model. This setup created a fixed inventory (no adjustments of supply were possible), a finite time horizon and a perishable product. As this new way of producing and selling fashion items became the norm in large segments of the fashion industry, pricing, and more specifically markdowns, could not anymore be decided by store managers based on rule-of-thumb and guesses. Markdowns needed to be optimized, leveraging the data generated by repeated short selling seasons.

The Gallego and van Ryzin 1994 paper used intensity control theory and applied it to the price-demand relationship, in order to control the intensity of the demand. The authors modelled demand as a Poisson process, which means that they considered the rate of arrival of customers in the store. Events (the arrival of customers) are unrelated to each other and independent of each other, they are random, but the average time between the events can be known. This rate of arrival (or intensity of demand) was modelled as function of price. The paper developed a closed form solution, providing the entire pattern of prices for a rate of arrival that is an exponential function. The model determined the intensity of demand that maximized revenue, as a function of the level of the inventory and the time left. The price decreases slightly after each sale, while a new sale slightly raises the price.

The model helped, in practice, to go beyond a crude trade-off: slashing the prices when one comes closer to the end of the season to get rid of the inventory left, but when making sales as time goes off, the product becomes scarce, and there is an incentive to raise the price. The model helped to move from that intuitive trade-off to an algorithm that leveraged the historical point-of-sale data, in order to control the intensity of demand by the prices, and maximize revenue much more efficiently than by setting a few classes in advance. The new model was

economics as an undergraduate. Garrett van Ryzin earned a PhD from MIT Operations Research Center in 1991, which was, and still is, a very multidisciplinary venture. The faculty members associated with the OR Center come from the economics department, various engineering departments, or the Sloan School. At the time of the publication of the paper, van Ryzin was an assistant professor at Columbia Business School, where he met Guillermo Gallego.

helpful because it was economical in its requirements for the estimation of the price-demand relationship: time series based on time of purchase data would do the trick.

This rested on some important assumptions. Demand for last year's casual tops should exhibit consistency with this year's. Discrepancies between the two years were treated as noise. The new dynamic pricing approach still retained a lot from the airlines revenue management approach. The customer was considered as requesting a specific product (a casual top, a specific fare for a specific flight) with a reservation price. The decision to optimize is whether to accept that request or not – or at what price in the new dynamic pricing model. The two authors did not consider that the customer was making a choice between alternatives. Moreover, the customer's reservation price, or willingness to pay, remained out of the picture.

The paper was important because it created a new class of problems, pricing problems under inventory constraints, that could be approached with dynamic programming, in an operational perspective. Yet the paper was immediately seen as important for two more specific reasons. First, it provided an entry point for revenue management into retail. Second, the airline revenue management problem could now be considered through the lens of the fully optimal solution, and this had an important effect on business practices.

Network Revenue Management: the Happy Marriage of Theory and Practice

The network problem was not invented by academic researchers, but came from businesses, in particular from hotels. In 1990s Marriott developed its own revenue management system, directly inspired by American Airlines.[28] The team was led by Marriott vice-president, Richard Hanks, and seconded by consultants from Aeronomics, a recently created company which sold revenue management systems to airlines.[29] The problem faced by hotel was so similar to the one faced by airlines that this translation is not surprising. In hotels too, capacity is fixed (the number of rooms in a hotel) and the product is perishable (the hotel room for a certain night, if not occupied that night, is lost; that is from the revenue perspective), there is a finite selling horizon.

However, the airlines model could not be straightforwardly translated to hotels. The room on a given night is not exactly the seat in the cabin of a scheduled flight from a revenue management perspective. In hotels, the length of stay is crucial. If the hotel is a business hotel located downtown in a convention city, Wednesdays could be in high demand, demand could even exceed capacity, while Tuesday and Thursday there might be empty rooms. A revenue management system optimizing revenue for each night, as in the classic airlines model, would accept a request for one night at a high rate on Wednesday, because it is a request for the high rate. Yet this sale could block the sale of a three nights stay from Tuesday to Thursday, the three days at the low rate. Depending on the prices of each of the two rates and the distribution of demand over the three days, it might be actually much more profitable to accept three days at the low rate.

[28] Hanks, R., Cross, R. and Noland., P. 1992. Discounting in the Hotel Industry: A New Approach. *Cornell Hotel and Restaurant Administration Quarterly* 33 (1): 15-23
[29] Aeronomics sent at Marriott Robert Cross, a figure of early revenue management systems, and the author of the most popular business book, which cover the early history of revenue management, including its development in hotels: Cross, R. 1997. *Revenue Management: Hard-Core Tactics for Market Domination*, New York: Broadway Books.

The classic airlines approach did not factor in the optimization what Hanks, Cross and Noland at Marriott believed was a good assumption on the demand: customers have a certain length of stay in mind when making a booking. This is the unit of demand, the one that should be forecasted and controlled for: not a room for a given night, but a room for a certain length of stay. The room rate can be set higher than if optimized for each night in isolation, accounting for the fact that the request is made for a specific length of stay. This made the optimization problem more complicated, since the control variable should be not only the room rate, but also the length of stay.

This problem was also almost immediately taken up by the airlines, still competing fiercely on their revenue management systems. The American Airlines first model and the EMSR improvement focused on a single leg, i.e. a single flight. Most airlines operated a hub-and-spoke system that was essential to their operations. Not many people want to fly from Fargo to Miami, but more people want to fly out of Fargo to somewhere. Therefore, Delta has a hub in Atlanta. The cheapest fare on a Delta flight from Atlanta to Miami could be set much higher than its competitors, as Delta attracts on that leg flows of customers from Seattle or Fargo, going through Atlanta to Miami.

The area became known as network revenue management, where the product sold consumes multiple resources on the network (multiple rooms for the length of stay, multiple flight legs in the airlines). New fare classes were created (in the airlines it was an origin-destination itinerary fare, called ODIF, with its own associated forecast; in hotels, various lengths of stay at a room rate). The data collected were the same kind of data, historical bookings data organized by fare class. The network problem increased the number of classes, from two classes for a single leg, to the collection of data on origin-destination-itinerary fare classes. The problem was becoming highly combinatorial (the so-called 'curse of dimensionality' in revenue management).

The Gallego and van Ryzin model provided a framework to deal with this network problem.[30] The network problem was one of allocating scarce resources (a limited inventory of seats or rooms) to competing uses (each fare class, each request for a specific length of stay at a specific room rate) which generates different revenue streams. This is the kind of problem linear programming is good at solving, if one considers for a moment that demand is known with certainty. In a linear program approach to the network problem, the objective function to be maximized was revenue (knowing the prices, the fares), demand is known for each fare (and one can list the resources composing each fare), and the constraint is capacity (of the plane), the variables are the allocation of seats to each ODF (numbers of seats).

The Gallego and van Ryzin approach had the following effects. First, this approach solved the network problem as it was framed in airlines and hotels, as a capacity allocation problem. Second, from there, the assumption on perfectly known future demand could be relaxed to account for the uncertainty on demand, which represented one more step towards the practical implementation of the optimal solution. Third, the linear program generated shadow prices, called 'bid prices' in airlines and hotel practices. This is a misnomer, as there is no buyer bidding. Bid prices are the shadow prices, reflecting the value created by relaxing at the margin the constraint on capacity. These could be interpreted as the marginal costs of capacity, or opportunity costs. Bid prices first provided a rigorous economic framework to think about the

[30] Gallego, G. and Van Ryzin, G. 1997. A Multiproduct Dynamic Pricing Problem and its Applications to Network Yield Management. *Operations Research*, 45(1), 24-41.

optimization problem. Second, from a practical point of view, shadow prices could be used as controls. Instead of storing numbers of seats available for each fare class, and then matching an incoming request with each number to make a decision, a burdensome procedure, bid prices can be used, i.e. one single number stored: all the fare classes lower that the bid price should close, all the fare classes higher than the bid price should be open. Interestingly, bid prices looked very much like dynamic prices.

The network problem tied academic research and business practices in the airlines and hotel industries. Yet two concluding remarks on the network episode could help to understand how, with the 1990s coming to an end, academic research on revenue management went ahead of business practices in these two industries.

First, the optimization: the airlines and hotels framed the problem as a quantity allocation problem, the allocations of resources (items in the inventory) to different fare classes, not as a pricing problem. The prices were not fully flexible, but set in advance, at the aggregate level, not through optimization. The allocation was static, made before the first request by a customer. The demand was considered to be exogeneous and independent from the allocation decision (which fare class were opened and closed). Changes in demand were only reflected through regularly actualized forecasts.

Second, the system: The network problem made the optimization more complex, and kept busy scores of operations researchers throughout the 1990s, but operated within the same system. The data collected were the same kind of data, historical bookings data organized by fare class. This was still the SABRE system, its data, its forecasts, and its decisions to accept or reject a new booking request.

The dynamic pricing approach solved problems and fostered a tight collaboration between research and business practices, yet it also offered a glimpse at the fully optimal solution, and that was ahead of the business practices in airlines and hotels in the 1990s.

Choice Models and the Legacy Systems

How limiting some assumptions made in practice were was first felt in the airlines context, by the same generation of academics.[31] The development of network revenue management in the airlines, mentioned above, made one of the basic assumptions of airlines revenue management, independent demands, increasingly weak. With the development of hub-and-spoke networks, the assumption of two independent streams of demand, deeply embedded in the booking systems, was becoming less and less accurate. Customers in the 1990s were obviously buying up (there is no leisure fare left, but I really need to go so I will buy a business fare), buying down (my company cut my expenses and the ticket is now too expensive, I will stay over for the weekend and buy the leisure fare). Customers could also consider different routes, non-stop flights and connecting flights, or could consider leaving earlier or later (in the early 1990s American had eight flights a day from New York to Tampa, plus connecting flights options). This seems pretty obvious, but accounting for this behavior in a revenue management system

[31] The story is told in a non-technical way in van Ryzin, G. 2005. Future of Revenue Management: Models of Demand. *Journal of Revenue and Pricing Management* 4 (2), 204–210. The most important technical paper in this literature is Talluri, K., van Ryzin, G. 2004. Revenue Management under a General Discrete Choice Model of Demand. *Management Science* 50, 15–33.

that automatized the decisions of opening up and closing down fares for all the flights was a different matter.[32]

The new models of demand developed for revenue management systems from the second half of the 1990s were developed for the airlines mostly. These choice models were imported from economics, and introduced in revenue management along the following lines.[33] The first step was to redefine the unit of demand. It was not a consumer that wants a certain flight at a certain price anymore. The aim became to model a consumer's choice at a specific point in time. In the second step, this choice was represented as a choice between a discrete set of alternatives. To each alternative a utility can be attached, which was a function of the attributes of each alternative. It was a regression model, with utility being a function of price, of some 'measures of convenience' (hours of deviation from the requested departure time), with some dummy variables (whether or not the flight is non-stop), and a noise term (to account for unobserved attributes). Third step, given the utility attached to each attribute, choice probabilities could be calculated. If multiplied by the number of customers making the choice, one has an estimate of demand.

Introducing this new way of modelling demand changed the optimization problem, and the literature came up with a new name to emphasize the difference: assortment optimization. Indeed, the aim was to use these new ways of modelling customers' choice in order to change attributes (the price, the time of the flight shown to the customers, the alternative flights shown as available to the customers), in order to see what the new demand would look like. The decision optimized in this new literature was not whether to accept or reject a booking request for a certain fare on a certain flight, given an allocation of seats made in advance to maximize revenue based on a forecast of two segmented streams of demand. The new decision was to decide which set of alternatives should be made available to customers, in order to shape their choices, and, ultimately, maximize the company's revenue.

Implementing choice models required a new type of data. In order to estimate the choice models, booking data for each fare (what was recorded in the SABRE system) were not enough. One needed to have access to data about choice outcomes: what alternatives the customers were presented with and which choice they made. Interestingly, operations researchers formulated these requirements as early as the early 2000s, while working within the airlines and hotels setups (with some projects on markdown optimization in retail, as mentioned above). Researchers had in mind data sources coming from loyalty programs. In airlines and hotels, loyalty programs could be used to tracks customers' behavior over time. Third party also sold more and more panel data. But in the literature, from the early 2000s, the proper data source for estimating choice models became web click-stream, which fulfilled all the requirements listed above.

---

[32] In the early 1990s, SABRE at American dealt with 5000 requests for booking per second at peak load time. At that time, the American SABRE system was the biggest database in the world, bigger than the U.S. government databases at the time.

[33] Choice models in revenue management were inspired by Luce: Luce, R. 1959. On the Possible Psychophysical Laws. *Psychological Review*, 66(2), 81-95, and subsequent developments in econometrics by McFadden. An important link between these seminal works and operations research is Ben-Akiva, M. and Lerman, S. 1985. Discrete Choice Analysis: Theory and Application to Travel Demand. Cambridge, MA: MIT Press.

From Demand Management to Marketplace Management (late 2000s to today)

[For now I am still collecting and processing the material on this part of the story. I just present briefly below a brief outline of the story for ride-sharing. There is much more to add, on ride-sharing, but also on online retailing and online advertising.]

Two things happened to revenue management in the late 2000s.

On the one hand, as explained above, research on models of dynamic pricing and on demand modelling were ahead of the revenue management systems used in airlines and hotels. In airlines and hotels, the implementation of dynamic pricing (making use of full price flexibility) and choice modelling (and the related assortment optimization approach) would require a complete change of the pricing system. The optimization model (the algorithm) should change, from static capacity allocation to dynamic pricing. The kind of data collected should change: from historical point-of-sale data to choice outcomes. The way these data were processed should change, from fare class demand forecasts to regressions and explorations in changing attributes to make predictions. Given the sunk costs and path dependency effects, given the deep embeddedness of revenue management systems within the operations of airlines, given how deep the fare class worldview was ingrained in airlines management and organization (with teams organized around the management of flights, not of markets), adoption of new revenue management systems in airlines and hotels was lagging.

On the other hand, in the late 2000s, early 2010s, companies in the emerging digital economy (online retailing, online advertising and ride-sharing) provided new fertile grounds for the development of revenue management. Firms run as marketplaces posed fine-grained, tactical problems of adjustment of demand with a less constrained supply, and more complex business objectives than straight revenue maximization. One can even go one step further and argue that in these industries, in online marketplaces, firms' operations did not just have a revenue management aspect. Revenue management was at the core of how these industries operated. One need to explain two dimensions: first, how revenue management moved into this new space, and second, how moving into this new space transformed revenue management.

At Uber and Lyft, from the mid-2010s onwards, the encounter of a new enabling technology (the three layers of mobile computing, GPS and maps) and a new business model (on-demand urban transportation) opened up a neat revenue management problem: how to manage demand when capacity is spatially limited (the drivers nearby)? Revenue management could leave the airlines, precisely at a juncture where airlines revenue management systems were evolving more slowly, and invest a new space.
Yet this first framing of the problem quickly evolved, and revenue management was transformed. First, capacity was less fixed than in the airlines or hotel setups, and can also be managed. Second, this in turn opened up a new spatial dimension for the optimization problem, as resources are not only manageable in time, but in space as well. Third, the so-called "capacity" or "resources" are actual drivers, not seats in an airplane. Their earnings are tied to the revenue management mechanism. This adds a third dimension to the problem of managing the marketplace, the problem of its stability.

Surge Pricing

Researchers from the operations research community were hired at Uber in 2015 to work on surge pricing.[34]

They first considered that problem as a classic revenue management problem. When a customer opens up the application on their phone, the number of drivers on the road is fixed. From a classical revenue management standpoint, improved by the dynamic pricing contributions made in the retail setup, the role of pricing should be not just to adjust demand to supply, but to maximize profit. A wealth of data on the customers' choice to book or not the ride after seeing the surge multiplier could be used as inputs to the model.

However, the specific features of the ride-sharing market forced the pricing team to tweak the classical revenue management framing.[35] The number of drivers on the road appeared to be more important that the straight maximization of revenue. Few drivers on the road mean a long ride to pick up the customer. This is a waste of time, both for the driver and for the platform. Many drivers concentrated in the same area result in many more rides. There is therefore an 'optimal number of idle drivers' at each point in time. Surge pricing aimed not just at dampening the demand and adjusting it to a fixed capacity, but also to manage that capacity, which is 'less fixed' than in the airlines context. Surge pricing aimed at incentivizing drivers, in order to get enough drivers on the road. Testing showed that there was no 'couch supply', i.e. drivers who were not driving and responding to a surge in price. However, surge prices incentivized drivers to stay on the road, move to a different area in the city (as the surges are distributed in space, not only in time), and on the longer run, to anticipate when and where to work.

Upfront Pricing

In 2016 Uber moved from a broker model to a merchant model. The broker model was based on a rate card that set a price for the time and distance of the ride. The platform then took a commission (usually around 20%) on the final ride for brokering the transaction. The shift to the merchant model meant that Uber would now buy the labour of the driver and then resell it, decoupling the cost (the payment made to the driver) from the price charged to the customer. The payment made to the driver was still based on a rate card that linked an amount of money to the time and the distance of the actual ride. But the price charged to the customer became set in advance, known as upfront pricing, and displayed in the application to the customer before they book the ride.

The motivation for this shift was first to show prices upfront to customers, and not surge multipliers, which can be scary (x1.3 of what?). The motivation was also that this change opened up the possibility to manage the directional flows through the network. A drive from downtown to far away in the suburbs implies that the 'resource', the driver, would waste time coming back, and therefore would be away from the network. Upfront pricing is a way to factor in this spatial externality, charging customers more for rides that undermine the maximization of the number of rides. This opened up a large computational problem, as origin-destination

[34] The team featured many OR figures, including some important contributors to the revenue management literature: Robert Phillips, Garrett van Ryzin from Columbia, Peter Frazier and Dawn Woodard from Cornell.

[35] As it is reported for instance in Lian, Z., van Ryzin, G. 2021. Optimal Growth in Two-sided Markets. *Management Science*, 67(11), 6862-6879; and in Yan, C., Zhu, H., Korolko, N. and Woodard, D. 2020. Dynamic Pricing and Matching in Ride-hailing Platforms. *Naval Research Logistics* 67(8), 705-724.

prices needed to be calculated in advance at any point in time, before the ride took place, not knowing how much exactly the ride would actually cost. The aim was to use upfront pricing to maintain an optimal number of idle drivers where customers are concentrated, in dense areas of demand, to maximize the number of rides there and dampen the demand elsewhere. The pricing mechanism supported the increase in the number of rides, on which the platform made small profits, in order to maximize the number of these rides and to support the growth of the market over the maximization of short term revenue.

Both surge pricing and upfront pricing were framed in a way that echoes many features of the classical revenue management problem. Fine-grained decisions on inventory control or capacity management (here drivers on the road) were integrated with decisions on pricing, in order to tactically maximize revenue. However, the ride-sharing problem presented three interesting new features. First, the optimization problem included the two sides of the electronic marketplace. Second, the optimization problem had a spatial component in addition to a time component. Third, profit making in ride-sharing was more tactical than blanket revenue maximization. In the ride-sharing setup, the resources are less fixed than in the airlines or hotels, and they can be managed through prices both in space and time. In turn, this makes the quest for revenue more tactical: decisions need to be made on where exactly (on which ride) to take profit.