Chapter 10
Operations Research in Rail Pricing and Revenue Management

Michael F. Gorman

10.1 Introduction

10.1.1 U.S. Freight Rail Pricing History

Market-based pricing decisions in U.S. freight rail are relatively new. Prior to 1980, U.S. freight rail prices were closely regulated by the federal government, leaving little flexibility for railroads to manage prices, and to a large degree handicapping the railroads financially. Since 1980, railroads have undergone an operating renaissance, rationalizing both physical plant and operating plan efficiencies. At the same time, while still obligated to "common carrier" commitments to continue service, railroads have been freed to price services according to market forces.

Interestingly, despite the Staggers Act's removal of rate ceilings, real rail rates have fallen over the last 40 years due to rail mergers, rationalization of rail capacity, and more efficient operations. (For more, see AAR 2013.) For many years after deregulation, railroads were massively over capacity with multiple parallel rail lines. Essentially, all incremental traffic that could be moved cost effectively was helpful to profitability. However, reducing parallel routes, falling real prices, and increasing traffic have created a rail network that is tightly constrained and congested. For this reason, careful pricing of existing and potential services is essential to modern rail performance.

M.F. Gorman (✉)
University of Dayton, School of Business,
Department of Operations Management, Dayton, OH, USA
e-mail: mgorman1@udayton.edu

10.1.2 Revenue Management for Rail: Importance

Freight rail revenue management is an analytical approach to pricing services that provides the maximum return from a fixed network of assets. In a sense, pricing action can be used to shape demand to better match the physical capacity of the rail network. In order to successfully implement a revenue management program, a railroad must understand both the behavior of its customers, and have a good definition of its capacity to provide service.

Because of the relatively high physical plant costs of rail (yards and rail lines), the economies of scale of individual trains (locomotives, crews) and relatively low marginal costs per load on an individual train (fuel and car), rail capacity is generally scarce even though the marginal cost per unit is relatively low. It is incumbent on railroads not only to allocate this capacity carefully so that service commitments can be met, but also that the available capacity be priced in a way that maximizes the return on investment in rail assets. In this environment, prices can be used to attract or deter incremental business, as deemed appropriate.

Because of the high fixed cost structure of rail, high volumes and high capacity utilization are mandatory for earning a reasonable financial return on the network. In this situation, the pricing decision must consider multiple capacity constraints and network effects of the pricing decisions of each service. Railroads revenue management decisions must sell the right combination of products, to the right customers, at the right time to maximize the return on the rail network. In some cases, incremental traffic can be very profitable to service as incremental traffic on existing trains, costing only additional fuel and equipment. The marginal costs of these loads are quite low, given the train is going to run in any case. These loads on average will lower the average the fixed costs of crew, line capacity consumption, and to some degree locomotives. In other cases, incremental traffic can be very expensive to service when the result is incremental trains which create demands for expensive locomotive, crew, line and yard assets, increasing average costs of service and reducing profitability.

10.1.3 Revenue Management for Rail: Challenges

Rail revenue management is similar to other industries with high fixed and low marginal costs, such as airlines, hotels, cruises, and the like. Given some relatively fixed capacity, appropriate pricing can maximize the return on that capacity. Unlike those industries, however, rail network capacity is considerably more flexible, both due to variable train sizes as well as the North American practice of highly flexible train services. Capacity is measured in a number of dimensions, such as car, locomotive, track, train, and yard. Further, that capacity is shared among competing business units. So, the unit of capacity is more difficult to define than, say, in the hotel industry where a room is a clearly defined unit of capacity.

As a result, pricing in the rail industry is very complex. First, the network is shared by numerous distinct traffic types: Intermodal, mixed merchandise, coal, grain, and automotive. Each has its own traffic patterns and seasonality, business shipping patterns and volumes. The vast majority of rail business moves on shipper-specific, medium-to-long-term contracts over which prices are predetermined. The structure of these contracts varies by traffic type; for example, intermodal tends to be shortest term, with a fair portion moving on spot market prices, and automotive and coal tends to be longest term. For this reason, for these lines of business, intermodal is perhaps the most conducive to revenue management techniques, followed by mixed merchandise. Coal, grain, and automotive have little short-term pricing flexibility, and capacities are often fully utilized in "unit" train operations.

To complicate matters further, many freight rail movements occur over multiple railroads, which then split the revenue according to predetermined agreements, limiting pricing flexibility and increasing revenue management complexity by requiring multiple railroads to agree on appropriate pricing.

Unlike airlines and hotels, railroad capacities are somewhat flexible, and fungible between product lines. In the hotel industry, an empty hotel room in one city cannot be substituted for a shortage in another; however, railroads can reallocate capacity between product lines. For example, an additional train could be created to provide service to excess demand in one product line and corridor, at the expense of capacity in another area of the rail network. Conversely, if demand is low, a train can be annulled, saving the capacity for other lines of business. In its simplest and most extreme form, railroads run trains only when they reach critical scale, known as a "tonnage" operation, and customer demand (not asset supply) determines the train operations. As such, a railroad's flexibility in service provision, traditionally viewed as an advantage to profitability and efficient operations, is a challenge to the use of traditional revenue management techniques which are generally based on the assumption of a fixed capacity.

Because of the complexity of pricing contracts, definition of a single unit of capacity, flexible operations and shared revenue between railroads, pricing decisions are extremely difficult. As a result, operations research models have not addressed revenue management in rail extensively; rather, such models have focused on asset capacity management given a level of demand, rather than transforming demand to better fit the capacity of the rail network.

10.1.4   Revenue Management for Rail: Recent Opportunities

More recently, as rail networks grow and become more complex due to mergers, ad hoc operations become progressively more challenging to manage. Downstream repercussions in the network from various on-the-fly service decisions can be unintended negative results. Further, customer service is more variable in a tonnage operation, making aggressive service-based pricing more difficult. As a result, more

railroads are tending toward running a (more) regular schedule with fewer extras and annulments, improving the feasibility of leveraging revenue management methods.

Railroads can establish revenue management strategies that balance the various needs of the different lines of business across the asset classes. Because the problem is relatively complex and ill-defined, most efforts in this area have been limited to one or a few dimensions on the problem, and few have been successfully implemented. Further, there are relatively few examples of advanced methods in revenue management in practice in U.S. freight rail. The remainder of the chapter describes progress and opportunities for revenue management in freight rail.

10.2   Analytical Techniques in Freight Revenue Management

Multiple analytical techniques have leveraged railroads' efforts to improve revenue management in the last 30 years since deregulation. In most cases, for simplicity, the pricing methods are applied to a single or a few dimensions of rail capacity. The review here focuses on US freight rail; for further reading in passenger and non-US applications, the interested reader is referred to Armstrong and Meissner (2010).

There are two major components of effective revenue management: Understanding customer behavior (level and price sensitivity of demand), and defining rail capacity (in a number of dimensions). The remainder of the chapter follows this organization; revenue management methods will be covered according to the following dimensions of the problem.

- Characterizing customer behavior: Estimating product demand.

  - Forecasting demand levels—prediction of traffic volumes and changes in demand due to exogenous factors.
  - Predicting customer price sensitivity—changes in volumes as a result of changes in price as a function of competitive factors.

- Characterization of rail capacity.

  - Train capacity—given train plan, selling available train capacity.
  - Equipment capacity—railcars and containers.

10.3   Characterizing Customer Behavior: Estimating Product Demand

A critical component of revenue management is understanding the underlying customer demand behavior. There are two primary measures of customer demand: Estimated response to price changes of the railroad, known as price elasticity, and any other factor that affects customer demand levels, such as seasonality, competitor

prices, changes to industrial production levels or international exchange rates given some price level, known as demand forecasting. Demand forecasts based on exogenous factors assume no pricing action on the part of the railroad. As price decisions are endogenous to the railroad's decision options, such changes are not subject to statistical methods and not included in firm-level demand forecasting. However, customer reactions to such price changes are not controlled by the railroad, and important to understand. Typically, but not always, statistical methods are used to estimate these customer behaviors.

There is a third component of customer behavior that has been studied as well, customer service elasticity. These lines of study look at the gain or loss of freight as the frequency, speed, and reliability of service improves or falls. With respect to traditional demand curve estimation, service level is a product attribute that can shift the demand curve just as a competitor's price change can. Because rail prices tend to be somewhat difficult to modify in the short run, an alternative is to modify service in order to attract or deter traffic. In so doing, costs and revenues are affected through service elasticities. The reader interested in a detailed review of service elasticity is referred to Small and Winston (1998).

10.3.1 Forecasting Demand Levels

Having anticipated customer demand as a function of exogenous factors and historical patterns is critical to revenue management. The demand forecast creates a baseline traffic level from which to plan capacity allocation and necessary pricing actions to affect change in anticipated traffic levels. As important as this demand forecast is, very little applied published research has been conducted on the subject in a practical, railroad-specific setting. Researchers have focused more on aggregate models of demand and modal split of freight flows at an aggregate level, which is not useful at a microeconomic level. Statistical methods, inventory-based methods, and utility maximization methods have been proposed and tested for projecting the total rail freight flows for a month or quarter based on the truck-rail modal split for broad geographic regions. However, such aggregate projections do not help the decision on how to price specific products, origins and destinations or days of the week. Thus, research to date on demand forecasting is of little assistance for revenue management. The interested reader is directed to Clark et al. (2005) for a survey of these methods.

Demand forecasting to support revenue management decisions is decidedly more practical in nature. A typical forecast will predict the number of shipments at varying degrees of specificity in geographical, commodity, and time period definition. For example, a forecast might project the number of shipments of a specific commodity or commodity group from an origin region to a destination region on a certain day, or even time of day. The level of specificity of a forecast is an important decision. The more specifically defined the demand, the more series there are to estimate, and generally the less accurately they are estimated. Often, the level of

specificity of the forecast relates directly to dimension of capacity of interest (e.g., car, yard, train). For example, to forecast equipment needs, the type of commodity must be specified in order to assess the type of car needed. However, for predicting locomotive needs, a railroad need only predict the total tonnage between yards, which is far less specific a demand forecast. In any case, to support precise revenue management decisions, a large number of demand series must be estimated.

In practice, demand forecasting is done statistically using traditional time series and regression-based forecasting methods. Of course, such statistical models at a detailed level of disaggregation are fraught with error. Railroads, unlike airlines and hotels, do not have reservations and thus have little forward knowledge of customer orders, so forecasts are more uncertain in freight rail. ("Car orders", or requests of shippers for an empty car to enable a future shipment is one forward indicator in advance of a shipment, but it does not require shipment on a certain date.) This deficiency of forward information has been another impingement on the progress of revenue management in freight rail. As a result, many pricing initiatives are considerably more aggregate and less precise in their efforts to define capacity and demand. For example, pricing may be applied at the train level to line capacity based on an average day, or price for a particular service may be based on quarterly demand.

10.3.2 Predicting Customer Price Sensitivity

The demand forecast becomes the baseline for estimating asset requirements to provide services. As the service requirements based on demand forecast are established, areas of capacity shortages and excess in the rail network can be identified. Pricing actions can incent or discourage additional shipments. But it is critical to estimate how much customer shipments will be affected by a change in price, which requires an estimation of customer price sensitivity, known as demand price elasticity.

Early efforts focused on statistical methods to estimate demand. For example, Winston (1983) describes demand and elasticity estimation for transportation demand. There are many other such works. While not explicitly related to revenue management, understanding transportation demand is a critical input to appropriate pricing. Most of these manuscripts use statistical methods such as regression to characterize demand.

Similar to demand forecasting, a significant challenge for estimating demand sensitivity is determining the level of fidelity to conduct a study. On one hand, customer-specific response to a specific origin–destination pair (and possibly day of week, time of day, and level of service) is of utmost interest, but such detailed data are sparse and unreliable. For this reason, many studies utilize more aggregate demand for which sufficient data density exists. For example, an elasticity study might look at monthly flows between regions. However, pricing action is often at a more specific product level.

Recently Gorman (2005) proposed a practical alternative based on an optimization-based alternative to statistical methods, based on optimization methodology, and observed profit margins in a market. The reasoning goes as follows: Railroads earn higher margins on products where demand is more inelastic, thus, the current margins are an indication of customers' price elasticity. This "implied elasticity of demand" has the advantage of having relatively low data requirements: only current price and current marginal cost. Railroads typically cannot attain a marginal cost, so an average variable cost is substituted. Gorman (2005) compares calculated implied elasticities with the intuition of market managers, and finds the estimated elasticities largely agree. The approach thus provides plausible elasticity estimates which are consistent with existing market prices. The method is dependent on some level of rational pricing; to the extent that current prices are severely suboptimal, the resulting elasticity estimates are biased. Further, the method also does not predict future levels of demand due to market forces; it is meant only to predict market response to changes in prices.

10.4   Research in Revenue Management Models

Given an estimate of shipper demand and its sensitivity, revenue management models attempt to maximize the return of some dimension of capacity. As noted, because of the uncertainty of demand and the flexibility of supply, relatively little work has been done in the area of revenue management in freight rail. What work has been done is primarily with respect to a single dimension of capacity to allocate, such as train space or container allocation. Some work has focused on service sensitivity of customers.

10.4.1   Train and Block-Based Capacity Approaches

The first freight rail-based researches to consider revenue management explicitly were Campbell and Morlok (1994) and Campbell (1996). Not surprisingly, this research converted the relatively successful approach of revenue management in the U.S. passenger airline industry and adjusted it to freight rail.

This research assumed a fixed train network with known capacities. The train network was assumed feasible with respect to other dimensions of capacity such as yard and line. Customer demand was assumed known and deterministic, and prices predetermined. As such, the revenue management model was based on deciding which set of customers who share these trains (i.e., intermodal, general carload) to provide service in order to maximize profits. The challenge of this model was to trade off various services, given the capacity each customer consumed on

a sequence of trains. As in the airlines, customers of different origin–destination pairs share capacity of intermediate trains; thus, the decision variable is in determining "block" capacity (allocation of train capacity by origin and destination of the block network) to be allocated amongst various origin–destination pairs. The model aims to maximize the expected profits of a set of blocks, given total block sizes do not exceed train capacity or customer demand levels. As in air revenue management, the complication arises in the complexity of defining a shared train network amongst blocks of different origins and destinations. In this model, capacity and block routings are fixed, and customer demand is accepted or rejected based on capacities. The work was never directly implemented by a major railroad.

However, in a simpler and more applied setting, CSX transportation experimented with a simple form of train-centric revenue management in the early 1990s. Several routes or trains were identified where the average tonnage and length of the trains were well below what could be handled with a standard locomotive consist. However, these trains were critical enough to the network that they were operated more than 5 days a week. The idea was to see if business could be attracted to these routes or trains by offering reduced rates to new customers who would move their freight only on these lanes.

To test this, the Operations Research group within the Service Design department developed flow maps of all the routes that regularly had at least 20 % available capacity using a standard locomotive consist. These maps were updated on a monthly basis and presented to the Sales and Marketing departments as well as the Finance department. In fact, there was a "Yield Management" team composed of members from Sales and Marketing, Service Design, and Finance that met each month to review progress and the most recent flow maps.

Recognizing that the key fixed costs of operation, crew and locomotive costs, would be incurred by these trains whether or not additional, incremental traffic was generated, Sales and Marketing teams were given reduced rates for these lanes that they could use to develop new business. New business was targeted because there was a goal to not diminish the revenue received from existing customers on these trains. No long-term contracts could be entered into with these new customers, since the space availability on the existing train service could not be guaranteed into the future. The team recognized that if the incremental business caused the need to add new train service, then the economic assumption that the fixed costs were being covered by existing business would no longer be valid.

This process was initially successful in generating some new business on these lanes. However, a problem was soon identified. Reports that were run each quarter to identify the profitability of customers used a process where the operation costs, including crew and locomotive costs, were allocated to *all* customers on a tonnage basis whose freight was handled on a given train. For the trains that were included in this initial Yield Management test, the impact of this cost allocation process was to make the existing customers look even more profitable since they were being allocated a smaller amount of the fixed operation costs than before. It also had the

impact that the profitability of the new customers who were brought in under the Yield Management pricing program looked abysmal since their rates were not designed to cover any of the fixed costs of the trains, and yet those costs were being included in this profitability analysis. The Yield Management team tried to have the algorithm by which costs were allocated changed, but without success. Given that salespeople's compensation and bonuses were impacted by these "profitability reports," sales that were based on the Yield Management pricing model soon dried up and the program was dismantled. Yield Management became linked to price cutting and unprofitable customers, and had a poor reputation within CSX for years to come.

More recently, a project with Amtrak used yield management techniques for the combined passenger and vehicle in Amtrak's "Auto Train" service product (Sibdari et al. 2008). This project is somewhat unique in freight rail both because of the joint passenger and vehicle decision, and because Amtrak's schedule and capacity is more fixed than the typical US freight railroad. They describe a discrete-time revenue management model for the single-leg Auto Train and evaluate three different heuristic solution methodologies. Reportedly, this approach is in use in the Amtrak revenue management department.

10.4.2 Service-Based Pricing Strategies

Kraft (1998, 2002) and Kraft et al. (2000) suggested another approach for rail revenue management. Railroads do not have fare classes as do the railroads, and prices are difficult to adjust; thus the approach has short comings because the rail industry is inherently different and more flexible than airlines in its capacity allocation. This line of research develops a multi-commodity network flow approach, where each shipment is a separate commodity. The model allocates potential demand over a number of different service options given a train network, maximizing expected revenues. Rather than allocating block capacity among customers, customers are assigned to different blocks based on their expected willingness to accept different service times. Critical to the approach is the assessment of the probability that a shipper will accept the service level of various routing options. As a result, demand is shaped by adjusting service levels in a way that is consistent with the train service network.

Other service-based models evaluate revenue implications for railroads. Strasser (1996) evaluates the development of a service-differentiated intermodal rail network and pricing impacts. This research suggests that service differentiation helps to enable revenue management strategies by allowing differentiated pricing by service type. Other thesis work (Nozick 1992; Kwon 1994) has considered service implications of network design and revenue implications. These projects are tangentially related to revenue management, but like the other service-centric projects described above, none are implemented in the U.S. rail industry.

10.4.3 Container-Centric Yield Management

In intermodal networks, container capacity availability and balance is a potential source for revenue management-based approaches. As discussed in the railcar management chapter of this book, empty repositioning is an integral part of service provision for railroads; cars and containers must be moved to where they are needed. Over the span of a month or a quarter, "strategic" container flows management must be balanced either through customer orders, or costly equipment repositioning. In the shorter term (e.g., 1–3 days), inventory of available containers in a geographic region must be "tactically" allocated profitably among tendered loads. Strategic and tactical container-centric revenue management approaches are described below.

From the railroad's perspective, some applied work has been published with container-centric revenue management objectives. Gorman (2001, 2002) discusses the use of pricing to help obviate such imbalances for BNSF intermodal. Instead of repositioning empties in an optimal way given an imbalance, pricing action can be taken to help balance the network in a profit-maximizing way. By raising prices in high demand lanes and lowering them in low volume lanes, imbalances can be reduced via pricing action. Gorman proposes a stochastic non-linear pricing optimization over a medium-term (e.g., quarterly) horizon. The work shows BNSF railway experienced an improvement of balance and therefore a reduction in repositioning costs while increasing expected revenues.

Since the early 2000s, the intermodal marketing company or IMC, which acts as a retail arm and third-party transportation management coordinator, has been taking ownership of its own containers, and thus has started to think about container-capacity based revenue management. Adelman (2007) evaluates strategic network pricing decisions in intermodal by modeling a dynamic fleet management problem on a closed logistics queueing network. Adelman's model leads to internal cost parameters similar to shadow prices based on network costs for improved dispatching decisions. The improved container allocation better balances supply and demand in the network. This work was not put to use in a practical setting.

In the short term horizon (e.g., 1 day to 1 week), container capacity in a geography in a geographic region is largely fixed because container repositioning takes considerable time. Gorman (2010) considers the decision facing Hub Group, an intermodal marketing company, when a shipper tenders an order. Given limited container capacity in the near term, Hub can accept the tendered order and its revenue, or reject the order in order to preserve the capability to accept a higher-revenue order that may be tendered subsequently. Further, the decision to accept an order should consider the anticipated future container supply and demand conditions at the destination of the shipment, which affects the future profitability of container capacity. This accept/reject decision does not allow pricing decisions, which are fixed in the short run, but manages container capacity in a way that maximizes expected revenue over the short run. Gorman suggests a simple probability-based heuristic based on expected revenues and the probability of running out of container capacity. Hub Group experienced both an increase in margin, a decrease in low-margin moves, and an increase in container velocity from container capacity-based load acceptance.

10.5   Future Directions and Opportunities for Revenue Management and Freight Rail

Recent research by Crevier et al. (2012) proposes joint capacity management and pricing decisions, attempting to bridge the gap between operations and pricing. This ambitious research expands the decisions beyond pricing given a capacity, and combines the two decisions. Working with Canadian National for practical input and data, they develop a largely theoretical approach that establishes both optimal pricing as train service provision, rail car handling, as well as capacities for handing railcars at classification yards. The ambitious project has not been implemented, but points in the direction of more integrated and holistic pricing and operations decisions.

References

Adelman D (2007) Dynamic bid prices in revenue management. Oper Res 55(4):647–661

American Association of Railroads (AAR) (2013) https://www.aar.org/keyissues/Documents/Background-Papers/A-Short-History-of-US-Freight.pdf

Armstrong A, Meissner J (2010) Railway revenue management: overview and models. Working Paper (available at http://www.meiss.com), Lancaster University Management School

Campbell K (1996) Booking and revenue management for rail intermodal services. Ph.D. thesis, University of Pennsylvania

Campbell K, Morlok EK (1994) Rail freight service flexibility and yield management. Proc Transp Res Forum 2:529–548

Clark C, Proulx B, Thoma P (2005) A survey of the freight transportation demand literature and a comparison of elasticity estimates. IWR Report 05-NETS-R-01. http://www.corpsnets.us/docs/other/05-nets-r-01.pdf

Crevier B, Cordeau JF, Savard G (2012) Integrated operations planning and revenue management for rail freight transportation. Transport Res Part B Meth 46(1):100–119

Gorman MF (2010) Hub group implements a suite of OR tools to improve operations. Interfaces 40(5):368–384

Gorman MF (2005) Estimation of an implied price elasticity of demand through current pricing practices. Appl Econ 37(9):1027–1035

Gorman MF (2002) Pricing and market mix optimization in freight transportation. J Trans Res Forum 56(1):135–148

Gorman MF (2001) Intermodal pricing model creates a network pricing perspective at BNSF. Interfaces 31(4):37–49

Kraft E (2002) Scheduling railway freight delivery appointments using a bid price approach. Transport Res Part A 36:145–165

Kraft ER, Srikar BN, Phillips RL (2000) Revenue management in railroad applications. Transport Q 54(1):157–176

Kraft ER (1998) A reservations-based railway network operations management system. Ph.D. thesis, University of Pennysylvania

Kwon OK (1994) Managing heterogeneous traffic on rail freight networks incorporating the logistics needs of market segments. Ph.D. thesis, Massachusetts Institute of Technology

Nozick LK (1992) A model of intermodal rail-truck service for operations management, investment planning, and costing. Ph.D. thesis, University of Pennsylvania

Sibdari S, Lin KY, Chellappan S (2008) Multiproduct revenue management: an empirical study of Auto train at Amtrak. J Rev Pric Manag 7(2):172–184

Small K, Winston C (1998) The demand for transportation: models and applications. Irvine Economics Paper, No.98-99-06

Strasser S (1996) The effect of yield management on railroads. Transport Q 50:831–844

Winston C (1983) The demand for freight transportation: models and applications. Transport Res Part A Gen 17(6):419–427