

Absolutely, Positively Operations Research: The Federal Express Story

RICHARD O. MASON

*Edwin L. Cox School of Business
Southern Methodist University
Dallas, Texas 75275*

JAMES L. MCKENNEY

*Harvard Business School
Soldiers Field
Boston, Massachusetts 02163*

WALTER CARLSON

*216 West Hilton Drive
Bolder Creek, California 95006*

DUNCAN COPELAND

*School of Business Administration
University of Western Ontario
London, Ontario N6A 3K7
Canada*

Federal Express Corporation has used operations research (OR) to help make its major business decisions since its overnight package delivery operations began in 1973. An early failure pointed out the need for scientific analysis. Subsequently, a successful origin-destination model followed by models to simulate operations, finances, engine use, personal assignments, and route structures influenced the conduct of business during periods of substantial growth. There were many false starts between the successes. CEO and founder Frederick W. Smith played a central role in the use of OR at the company: he established a relationship with OR and management science personnel and this relationship supported the growth and success of the company.

This company is nothing short of being the logistics arm of a whole new society that is building up in our economy—a society that isn't built around automobile and steel production, but that is built up instead around service industries and high technology endeavors in electronics and optics and medical science. It is the movement of these support items that Federal Express is all about.

—*Frederick W. Smith*

Frederick W. (Fred) Smith developed a vision of a business, brought together the key people who would make up his ini-

tial management team, and scraped together enough funds to buy the first aircraft for his fleet: French-made Dassault

Falcon-20 twin-engine executive jets.

His idea was to provide overnight delivery of small, high-value items, such as pharmaceuticals, aerospace components, and computer parts. An avid pilot, Smith had chosen the Falcons because they best satisfied the constraints the business had to work under. To be successful, the business had to be free to change its routes and schedules frequently and readily. Civil Aeronautics Board (CAB) regulations on air cargo at the time, however, did not permit this degree of flexibility. Consequently, the

Turboprop planes were too slow.

company was forced to operate as an air taxi service and use aircraft that had a payload of less than 7,500 lbs. Early studies had shown that turboprop planes were too slow to meet tight overnight deadlines. Executive jets were the only aircraft that would. Compared with other executive jets, the Falcons were sturdier, had lower operating costs, and featured a GE engine that was more powerful and considered more reliable than the Pratt and Whitney or Air Research engines. Equally important, Dassault had 22 new Falcons sitting in the New Mexico desert and was willing to deal. With a 56-foot fuselage and 50-foot wing span, the Falcons had a range of about 1,500 to 2,000 miles flying at about 540 miles per hour and, when appropriately modified, could carry a payload of about three tons. Given Smith's requirements, the Falcons were about the only satisfactory aircraft—"the best of the worst," as some one put it—for an overnight service pick-

ing up and delivering small packages.

The founders were familiar with the South, and this seemed an appropriate place to begin their new business. They chose 11 cities, mostly located in the South and Southeast, to make up the first route network, which they would operate in a hub-and-spoke manner. They would fly outbound packages into Memphis, sort them, and then forward them to their inbound delivery destinations. FedEx thinks of things from the point of view of the package. So, in FedEx parlance, *outbound* refers to traffic moving from the station of origin to the hub (or towards its destination); *inbound* from the hub (or any intermediate spot) to the destination.

Initial sales calls had generated a lot of enthusiasm and large estimates of volume. Expectations were running high. Smith even worried that his tiny fleet would be overbooked. March 12, 1973 was set for the inaugural service. That evening Smith, his founding team, and a few of his major investors waited anxiously as the Falcons descended into Memphis with the first day's load.

Henry Meers, an investment banker, recalls the evening's events. "I saw the anguish on their faces as they waited. Most were very worried about their future. It was a critical moment for all of them as they finally crowded around the Falcons and the cargo doors opened. But as they gazed inside, there was bitter disappointment. There were only six packages and one of them was a birthday present Fred Smith sent to his close aide Irby Tedder" [Sigafos and Easson 1988, p. 59]. The next two evenings showed little improvement, and after only three days of operation, the

company discontinued air delivery service. To maintain its customers, the company found alternative ways to transport their packages while it took stock of its situation.

The First Model

After getting over the initial shock, Fred Smith appointed a task force headed by Charles Brandon, a physicist and an ardent flyer he had met in 1962, when both were home from college. When Smith was contemplating his new business in 1972, he asked Brandon, "Could you use a computer to better schedule the aircraft?" "Definitely," was Bronson's reply. Brandon soon thereafter became an advisor to FedEx and an early employee. Smith's charge to Brandon was straightforward: Where have we failed? What should we do now?

The team members closeted themselves in a conference room at Little Rock's Worthen Bank to begin their soul-searching, working 18 to 20 hours a day for 15 days straight. Their approach was classic. They began by challenging the assumption that the cities the founders knew best were best for doing business. This led them to focus on what characteristics made a city a good candidate for their route structure, and they concluded that they should include some cities because they generated a large volume of outbound small packages; others because they received a lot of small packages. Moreover, the final set of cities should account for a substantial amount of the total volume within the system.

Given these guidelines, they next acted on their intuitive ideas in a systematic way to find data or indicators they could use to select cities and to decide on a route system. First, they formulated the problem in terms of an origin-destination flow model.

They evaluated approximately 112 cities (yielding 12,422 cells = $112 \times 112 - 112$) to determine which cities they should include in the system. They developed data on origins from statistics brought by S. Tucker Taylor, a former consultant to Smith and member of the first model team, from the Civil Aeronautics Board on enplanement volumes and from other sources on such factors as population, employment, and business activity by SIC codes. They adjusted gross enplanement data for each of the 112 origins by subtracting out heavy freight traffic and then further massaged it on the basis of assumptions they made using SIC Codes (revealing types of businesses and, therefore, indicating potential for small package traffic), size, weight, and so forth. They used all of this to establish a coefficient of outbound market potential. Since no destination data was available, they assumed that a city's inbound market potential was proportional to its percentage of the total population of the 112 cities. They multiplied the outbound potential coefficient for the first city (the origin) by the inbound coefficient for each second city (the destination) to obtain a raw estimate of the one-way total volume for each traffic lane. They hung huge sheets of butcher paper on the walls and, with the aid of a hand-held HP-35 calculator, one by one filled out the cells of the matrix by hand.

Next the team went through the *Official Airline Guide* and the *Air Cargo Guide* to determine the availability of day and night air cargo flights, airline belly cargo, and other competitive means of transportation. They devised a simple set of rules and used them to adjust the raw origin-destination figures to obtain a FedEx expected "share"

of the predicted volume. Again they entered the figures by hand.

The resulting matrix was extremely revealing. It showed that the types of businesses a city had and the nature of its air cargo competition could affect its attractiveness for FedEx. Rochester, New York, was an example of a city with a quite vibrant economy—Kodak, Xerox, and other volume shippers of small parts were located there—that had poor air cargo service. It proved to be an inviting, economically viable location. The analysis showed that several other cities had good potential and poor existing service and that FedEx could expect to get about 30 percent of their small package business. Some cities in the original set of 11, like Jackson, Mississippi and New Orleans, paled in comparison. Delta Airlines, which carried some small packages on passenger flights, provided good service to these cities; consequently, FedEx could expect to capture only about five percent of their market because they lacked small package volume. As a result of these analyses, they combined the newly identified cities with a few major market centers, such as New York City and Chicago, to define a 26-city system, one that was quite different from the original 11-city system.

On April 17, 1973, a rejuvenated FedEx began serving its 26-city system. In one of the great reversals of fortunes in business history, the new plan of operations worked. Twenty-three years later, according to its home page (www.fedex.com), "FedEx is the world's largest express transportation company, providing fast and reliable services for important documents, packages, and freight." It delivers more than two mil-

lion items to over 200 countries each business day. FedEx employs more than 110,000 people worldwide and operates 500 aircraft and more than 35,000 vehicles in its integrated system.

The first model and the planning system that evolved around it have served as the foundation for the \$8 billion plus corporation that is the Federal Express Corporation today. They also established corporate values for analysis and modeling that have become integral to the way FedEx does business. In the process, founder and CEO Fred Smith has become a dedicated user and supporter of operations research and in a real way an active participant in it.

Come FLY with Me

A delicate balancing act is embedded in Smith's business concept. On the one hand, the company must find the most lucrative cities with which to do business. On the other hand, it must pick up, transport, and deliver packages from and to these cities efficiently under the constraints imposed by overnight delivery. Soon after operations began, the need to address the other side of the equation was apparent.

Norman B. Waite, a consultant with a degree in mathematics from Johns Hopkins, joined Charles Brandon to address this problem. The result was FLY, a deterministic model based on known parameters that simulated the operation of aircraft and facilitated the construction of flight schedules. The model drew on a database that contained information on the characteristics of airports throughout the US, the operating characteristics of the aircraft, loading times, unloading times, prevailing winds, flying times, and such other factors as air traffic control delay times, taxi delay times,

and instrument approach times. It simulated the flight of a given aircraft comprehensively from gate to gate. They used the FLY model to calculate such factors as flight times and engine cycles. They then used the results to estimate crew requirements, maintenance schedules, brake component repair schedules, and other operational demands. A few months after the first successful flight in 1973, Federal Express was renting time on an IBM 360-67 at National CSS Time Share, and Waite wrote the original code for FLY for the 360 in PL/1.

Shortly after FLY became operational, Brandon used the model to evaluate FedEx's expansion possibilities. Smith planned to serve as many as 82 cities with 33 Falcons. The expansion plan, however, required additional capital. The \$10 million seed money Smith had begun with was running out. Meanwhile, the second round of financing was in jeopardy. FedEx had to find new investors. But it had to overcome two major hurdles to get them to sign up. First, most potential investors had difficulty with the hub concept. ("Do you mean that a package sent from New York City to Newark flies all the way to Memphis and back?") Second, they wanted some proof that they would get an adequate return on their investments. Using FLY, Brandon constructed a model of a proposed new structure featuring 82 airports served by 33 Falcons. In a barnstorming tour, Smith, Art Bass (who became president and chief operating officer in 1975), and others used the 82 airport model to convince investors and employees alike that the company's basic concept was sound and that the new system was economically attractive. With the

aid of the model, the team obtained the cash it needed and added the new cities and aircraft to the system. At the time it was the largest venture-capital start-up financing in US history.

With FLY, the team could ask what-if questions and test its ideas. In one influential test, Brandon, as director of management information systems, demonstrated that Los Angeles was a "magic" city, a prime place for the next expansion. The company's resources were tight, and Smith posed the question, "If we could obtain just one new plane, what city should we add?" They used FLY to perform what economists call a marginal analysis, and it showed that Los Angeles drew and pumped substantial volume from and to other cities already in the system, creating economies of scale. Los Angeles had five to 10 times more positive impact on FedEx than the next highest city. FedEx made plans to include LA in its route structure.

The Three-Model Planning System

With these successes, modeling became a way of life for FedEx. Soon it developed a three-model management planning system:

- (1) It used an improved origin-destination flow model to determine the what, when, and where of package volumes from and to actual and potential cities in the system.
- (2) It used FLY to produce schedules and determine resource requirements for these cities. It also ran FLY using actual past volumes to review performance, to test other options, and to recalibrate its coefficients.
- (3) FedEx created a financial planning model to show the overall economic and financial implications of alternative route structures and flying schedules.

It used these three models in concert to make ongoing operational decisions as well as crucial strategic decisions. During the fall of 1973, for example, the three-model system more than paid for itself. Heretofore the price of fuel oil ranged between \$2.75 and \$3.50 per barrel. But, in October, the Arab members of OPEC announced a drastic reduction in its production of crude oil and an embargo directed at the United States and the Netherlands. World crude prices nearly quadrupled by the end of the year. Fuel was rationed in the US, and it appeared that FedEx would soon be driven out of business. Convinced that he needed to protect his company from impending disaster, Smith headed for Washington, DC—but not before requesting the OR team to run its models based on different assumptions of fuel oil prices and availability.

Brandon, with the help of Bill Arthur, estimated the growth in takeoffs, landings, and average flight length for the next five-year period and fed this data into the models to determine fuel consumption and costs. They added an optimistic fudge factor, just in case. J. Tucker Morse, FedEx's internal company counsel, immediately relayed the results to Smith in Washington. Smith's first visit was to Representative Wilbur Mills, who, after seeing the results, called Secretary of Energy John Sawhill. Sawhill reviewed the computer estimates and approved a special allocation for FedEx, enough to keep the 82-city network operating for several years. As it turned out, because of the team's generous assumptions on consumption, FedEx received a greater fuel allocation than it actually needed during the panic.

Economics and Modeling Systems

During the early 1970s, the underlying economics of FedEx's business became clearer. FedEx uses high-cost equipment and facilities intensively for a short period of time each night. Efficient asset utilization is made even more difficult by a set of precedence conditions, such as "the first plane cannot leave until the last package is sorted." Because the system's activities are interdependent, speeding up one activity does not necessarily improve anything else. The only economical solution is for each activity to carry as heavy a load as possible, but not so heavy as to exceed the capacity of the assets allocated to that activity. Consequently, capacity planning is essential, and operations must be conducted with a very high load factor.

It is challenging enough to develop one high-load-factor schedule for a stable system, but the problem Brandon and his crews faced was even more daunting. Because FedEx was adding customers, cities, and aircraft constantly, it needed to radically reschedule its routes at least once every month, while simultaneously insuring that every flight had enough capacity to keep the system from blowing out and leaving packages on the ground. Many common carriers operate at load factors hovering between 60 and 69 percent. With the aid of the three-model system, FedEx was generally able to maintain a load factor between 82 and 93 percent (its average was 87 percent) while reconfiguring its route structure every month, implementing its schedule changes within a few days, and keeping its financial performance in line. This was an accomplishment. As late as 1991, another major air freight carrier took

over a month to do the same calculations and to reschedule its operations.

Building the OR Team: Ponder and Hinson

Operations research was to be central and ongoing at FedEx. Brandon sought help and in 1974 hired Dr. Ron Ponder of Memphis State as a consultant. Ponder's student, Joe Hinson, also worked as a consultant during the summer of 1974. The following August, Hinson became a full-time employee. Ponder joined FedEx as director of OR in 1977.

In addition to his skills in applying OR, using GPSS tools for model building, and his expertise in database technology, Ponder proved to be a very capable manager. He built an OR team of 25 that became known as one of the finest OR groups in the world. When Ponder became a senior vice-president in 1980, Hinson became managing director of OR. He now reports directly to Fred Smith.

Joe/Engine

The first project Joe Hinson tackled in 1974 was scheduling maintenance on the General Electric CF-700 turbofan engines used in the Falcons. Since the engine had been used primarily in executive jets, it had not logged the hours or cycles necessary to get FAA approval for extended intervals between maintenance sessions, a factor known as time between overhauls (TBO). Without a large statistical database from which to estimate reliability, the FAA set FedEx's TBO on the Falcons at a stringent 600 hours. Overhauls took about two to three weeks, and since each plane flew about six hours a day, overhauls were required every 100 days. This was a very tight constraint for FedEx's small fleet.

Joe Hinson built a model to develop an acceptable overhaul schedule. Officially called "Engine," it was fondly known throughout FedEx as "Joe/Engine," because of the computer file name that appeared at the top of every printout. Its results were alarming. Given FedEx's expanding schedule, the model revealed that the required removal rate would soon surpass the company's total fleet capacity and bring operations to a halt.

Further events deepened the crisis. The FAA and GE notified FedEx's maintenance and engineering (M&E) department that GE had used a defective die to cast a critical part of the engine called the "bucket." As a result, this rotating blade tended to fail more frequently than expected. It appeared as though a bucket might fail every 50 cycles of use. A cycle is completed every time an engine is started, idled, powered up, idled, and then shut down. FedEx's exposure was much greater than most other users because its planes took off and landed so many times. A typical FedEx plane consumed several cycles a day, implying that every plane had to be taken out of service at least once a month. Worse, a bucket failure often destroyed the entire engine.

The model showed that there were not enough engines in the world to fly FedEx's fleet hours given its forecasted continual need to replace buckets. The OR team, working closely with the M&E Department, used the model to manage a pool of engines (FedEx rented some additional engines) and averted the crisis by developing strategies for swapping engines from aircraft to aircraft and by dynamically re-scheduling overhauls. With Joe/Engine,

FedEx kept its fleet in service until GE developed an improved blucket, and the crisis was averted. Subsequently, FedEx set up an engine reliability program and compiled a database of FedEx's actual performance. Based on this new evidence, the FAA eventually approved a 3,000 hour TBO for FedEx's CF-700 engines.

How about Pittsburgh?

Hinson determined that FedEx needed a forecasting model to bridge the gap between the origin-destination model and FLY and to sharpen its estimates of daily package counts for each city. Although FedEx had only about eight months of actual operations data at that time, Hinson acquired Air Transport Association's gross data and devised a way to create a seasonal index and to estimate growth factors. Using a crude form of exponential smoothing, he introduced each night's actual counts into the model every few days or so to keep its coefficients current. Early in 1975, he translated the FLY model for the newly acquired Burroughs B-6700 computer and programmed the forecasting model to accompany it.

FLY, coupled with the forecasting model, became a kind of early warning system. In March, the models predicted that with a fleet of only 33 Falcons sometime during the fall of 1975 FedEx would run out of lift, that is, its package volume would exceed its total aircraft capacity. Despite the fact that FedEx was barely eking out a profit, it had to do something.

Mike Staunton (then head of scheduling, now vice-president for global operations, scheduling, and control) suggested that a bypass hub in Pittsburgh might reduce some of the pressure. In the transportation

business, a bypass hub receives, sorts, and redistributes packages without forwarding them to the central hub. That is, it operates in parallel with the central hub and thus functions differently from a set of hubs operating in series in which packages are moved from hub to hub. Working together, Mike and Joe ran various simulations using the models and determined that a mini-hub in Pittsburgh would indeed be economical, but only if FedEx could use DC-3s for supplemental lift. The company had no DC-3s at the time. Consequently, it let contracts with local DC-3 operators and soon became the largest charter operator of DC-3s in the US. Pittsburgh was the first multiple hub to be run in parallel with Memphis, establishing an approach to multi-hubbing that is still in use for subhubs. Significantly, as FLY and related models have been expanded, the OR department has used them to evaluate many other proposals for additional subhubs.

Autoroute

The Pittsburgh solution proved to be temporary. FedEx faced an even bigger crisis, almost immediately. Late in 1975, the company would need planes with a much greater payload than the Falcons if it was going to handle anticipated increases in volume and earn an acceptable profit, but CAB regulations still prohibited companies like FedEx from using them. CAB consequently turned down FedEx's initial request to buy larger DC-9's.

The aircraft capacity problem went to the top of Smith's agenda and became the subject of almost daily conversations between him, his assistant Tex Weise, Staunton, and Brandon (then vice-president for operations planning). Reporting to him were airline

scheduling, headed by Staunton; telecommunications, headed by Jack Cockrill; and OR, headed by Ponder. He asked Hinson and three other OR personnel to study the capacity issue.

As has become the custom at FedEx, OR was represented in the weekly senior management meetings. During this era, Brandon attended the meetings, often accompanied by Ponder or Hinson or some of the other members of the OR team. This special role was accorded the OR group because Smith, who has a degree in economics from Yale and is a liberal arts and military science buff, had learned about how OR worked, was impressed with the results, and had come to trust the advice he received from the OR team. The payload problem was so hot that Smith met with the OR group often for extensive discussions between the regular executive sessions.

The team had originally built Autoroute as a research tool to help people understand the underlying structure of the aircraft scheduling problem. Joe Hinson believed that it would provide valuable background knowledge for solving some of the company's operations planning problems, and Brandon had encouraged him to explore the possibilities. Now, they used the model to determine the extent of the impending crisis.

Autoroute is a heuristic optimizing, scheduling model that takes loads, cities served, and constraints, such as the latest hour an aircraft can depart a city and the earliest it can arrive at its destination, and generates a complete single-hub-system flight schedule. Over several months, Hinson and his OR team made repeated

runs of the model, making different assumptions about the composition of the fleet, usually by substituting combinations of two or three different types of aircraft at a time. No matter what combinations they used, the result was the same: FedEx could not remain profitable unless it was allowed to use larger planes.

Smith was troubled and moved into action. He took an apartment in Georgetown, leaving COO Art Bass to run operations in Memphis, and began a series of wide-ranging lobbying rounds in Washington, DC. Virtually no one connected with airline or air cargo regulation escaped his attention. This effort took nearly a year. Armed with the Autoroute results, Smith became a strong advocate for an air cargo reform bill. In March 1977, he testified before the senate aviation subcommittee and asked CAB for a certificate of public convenience and necessity. He appeared before the subcommittee again in August, testifying in support of regulatory reform. This model-based lobbying eventually got results. In November, the House of Representatives passed the Air Cargo Reform Bill, and soon thereafter President Carter signed it into law. FedEx almost immediately filed for permission to fly larger aircraft. It used Autoroute to evaluate the use of Boeing 727s, among others, on FedEx's high volume routes, and in December 1977, placed an order for seven 727s.

Human Resource Planning

The 727s brought about new problems with personnel planning. As FedEx expanded, introduced more hubs, and launched a more diverse fleet, it had more difficulty making sure its jobs and pay scales were in order. Of special concern

were its pilots, who had to be qualified on each new aircraft and who had many job opportunities elsewhere. The pilots considered unionizing, and the company wanted to show them what they could reasonably expect in career earnings and types of aircraft to fly.

The problem was formulated as a multi-period, multi-stage stochastic (Markov chain) transition model. The OR team used fleet and operating requirements for a decade or more in the future to lay out pilot requirements by aircraft type for each period. The characteristics of the existing crew complement served as the starting point. The model then simulated the movement of each pool of pilots in a given class through these periods. Statistical data was used to determine the probability that a pilot would fail a medical examination, fail to upgrade, accept a job with another company, and other events that could affect his transition from one state to another.

The model calculated the pay for pilots, first officers, and others and was used to show them what their expected career path at FedEx would look like. Since the company was expanding, the future was promising for most pilots. This result proved very useful in labor-management negotiations. FedEx also used the model to forecast gaps between the supply and demand of personnel and to indicate what its future hiring and training needs were. This human resource planning became an integral part of the company's overall planning process.

Perhaps We Need a "Wizard!"

As FedEx's volumes grew, it became increasingly difficult to handle requests and inquiries from shippers. Until late in 1977,

each city station bought or leased its own equipment and ran its own phone system. Moreover, station managers jealously guarded their control over the process. "Our customers," a typical comment went, "want to talk directly to their favorite agent right here at home." Unfortunately, there was little call-answering discipline and there were no established standards. Vital information was constantly being lost. In short, it was chaotic. One day Tucker Taylor walked into Charles Brandon's office and asked "Why can't we have our own Wizard of Avis?" The "Wizard of Avis" was a computerized reservation system for car rentals that featured a centralized telephone system for accepting and routing all reservations throughout Avis's system. The possibility of a wizard seemed to Brandon like a question that OR could answer, and he moved quickly to propose a project to Smith. At first he ran into a stone wall. There was a sharp division of opinion within the company on this matter, and only after much heated discussion did Smith approve the development of a prototype dubbed Project Sydney (for *Sydney* Tucker Taylor). FedEx formed a project team consisting of representatives from EDP, telecommunications, field operations, sales and marketing, and OR. OR began the initial modeling, which would serve as the foundation for building the system.

Some preliminary model results suggested that a central telephone answering system was the best solution. A prototype for a centralized system was programmed for the Burroughs B6700. In late 1977, FedEx conducted a test in Newark, at the time the worst-performing city in the system on the basis of company statistics and

employee performance, and in Memphis. The prototype system, though limited in capacity and reliability, was a vast improvement over the existing system. "For the first time, the people in the stations could do their own work," one executive observed, "because they were relieved of answering the phones and processing the orders." Customers were also served better. Other cities immediately demanded that they too be added to the system. Despite some known flaws, FedEx rolled the prototype out to 10 or 11 other cities. Ultimately FedEx handled all customers' incoming calls in centralized call centers, which then sent out requests to dispatch centers located in each city FedEx served.

Given the stunning success of Project Sydney, Smith approved a project to build a full-scale production real-time system—eventually called Cosmos—to serve the entire FedEx system. At the time, there were many unresolved questions concerning how many call centers to establish, where to place them, and how to staff them.

Hinson and Ponder, with the aid of OR analyst Michael Sternad, built a customer service system model to explore the possibilities. The model clearly showed that length of call was the most important factor in determining call-handling performance.

Queuing analysis further revealed that it was quite easy to overload the system and to bring it down. In fact, a sensitivity analysis applied to the model, revealed that the length of the call was the main cost driver for the entire system. Slight decreases in the average length of call resulted in substantial cost savings; slight increases resulted in a substantial cost increases. Further studies suggested other methods for

reducing the time each call took, such as improving the agents' answering behavior, routing calls more efficiently across time zones, and modifying other procedures. FedEx used these results to develop an overall business plan for both the telephone system and the computer systems and to determine the optimum locations for the call centers. The call center study results served as the fundament for the subsequent development of Cosmos.

Ultimately, FedEx decided to build more call centers than the model indicated because its executives believed that the additional centers provided more flexibility, avoided possible regulatory issues, and would serve as a positive labor relations move—all factors not fully accounted for in the original formulation, but which emerged during intense discussions with Smith and others. Such give-and-take discussions, in which model results are pitted against executive understanding, have become common practice at FedEx.

Hubba, Hubba

Operations research influenced one decision that had an enormous impact on FedEx and set its course for over 15 years. This hotly debated decision was to build a SuperHub, and as Charles Brandon observed, "OR was applied effectively here and *reversed* the four-hub serial design decision in favor of a single SuperHub. It saved the company!"

When the network was small, its volumes low, and the Falcons the only aircraft, a single hub in Memphis was clearly the best choice (although the Pittsburgh experience had shown the value of parallel hubs). However, the dramatic growth of the business and the advent of the 727s—together

with the possibility of even larger aircraft in the future—raised questions as to whether this strategy was still optimal. Indeed, several key executives were convinced that a series, multihub system—most favored four—was the way to go.

Four hubs appeared to be a done deal.

They reasoned that FedEx could use small aircraft to carry packages to each regional hub and that it could use larger planes to haul packages from hub to hub. A great deal of momentum and enthusiasm built up within the company as people maneuvered for a four-hub network and, as one executive put it, “Four hubs appeared to be a done deal.”

Hinson, Ponder, and Brandon were not convinced. OR had not yet studied the problem in any depth. And so, as the front office moved toward a four-hub system, they began to model the situation. Their first task was to expand their modeling approach. Autoroute, it turned out, could be used to estimate the performance of feeders into a regional hub, but it was incapable of determining whether or not a given hub should remain in the final system network. Since the final solution must contain a discrete number of hubs, they needed an integer linear programming model (ILP). Consequently, they constructed a two-phase model consisting of Autoroute to perform within hub analyses and an ILP to choose from among the candidate hubs. They made exploratory runs assuming between one and eight hubs and also addressed fleet size and composition issues.

It was standard practice to discuss even very preliminary OR results in senior management committee meetings—Fred Smith was an especially eager participant. This dialogue served to change assumptions in the model and to suggest other approaches.

The first model runs indicated that a single hub was still the preferred solution. Every time this solution was presented, however, a heated debate ensued. Smith and others continually came up with new ideas about the cost of aircraft, the nature of the facilities used, and the like. With these new considerations in hand, the OR team would head back to its quarters to modify the model and to run it a few more times. This continual dialogue, pitting executive assumptions against their implications as derived by a model, lasted for about a year. Throughout this time, the single hub solution generally dominated in the model results but not, it turned out, in many executives’ minds.

In the end, Smith decided on the single-hub system—now dubbed the SuperHub. The model revealed the underlying economics clearly. With high cost aircraft and facilities used intensively for compressed periods of time under conditions that demanded near perfect performance—that is, failure to service a plane or to pick up a batch of packages could not be tolerated—the single hub worked best. Given the technology and volumes at the time, the SuperHub design was more flexible; it required only one big sort, reducing total sorting time, and it made crew domiciling easier. The four-hub approach had three major flaws: it caused low utilization of high cost assets and, hence, had a poor ROI (in some cases negative), it had a low toler-

ance for faults because it did not have the built-in redundancy that the single hub did, and it consumed almost all of the 12 hours FedEx had available to process work to meet its commitments. Hence, it did not allow much slack nor did it provide much room for growth. Memphis remained the single hub in the system, and in July 1979, FedEx broke ground for the massive, highly automated SuperHub.

Soon after the SuperHub was opened, the model revealed additional advantages to the concept. As a competitive move, Smith wanted to change the committed delivery time from 12:00 noon to 10:30 am. The model showed that the FedEx system could deliver on that promise. Furthermore, a competitive analysis application of the model revealed that UPS, FedEx's major competitor, could not effectively meet this deadline. In late 1982, Tom Oliver, head of FedEx's marketing, announced, "We will offer 10:30 delivery, more service options, Saturday pickups, package tracing, and call-backs to shippers informing them that the packages have been delivered," and he did this knowing that his company could deliver on this promise and that, for the foreseeable future, UPS could not. All of this was consistent with Smith's strategic emphasis on quality and reliability of performance.

The technological infrastructure of the company was now pretty much in place: a mixed fleet of small planes providing service to the smaller stations and larger aircraft on the higher volume runs funneling through a SuperHub plus centralized call centers for answering incoming phone calls, which were in turn dispatched to local stations for response. This infrastruc-

ture, which was predicated on OR results, called for a much different approach to FedEx's organizational design. Smith understood this, and in April 1978, he reorganized the company from its previous decentralized form into a highly centralized company, providing as an explanation that "various approaches were modeled against our present and future tasks."

Tracking toward the COSMOS

As volume growth accelerated, the problem of making the stations work better and servicing customers quicker and in a more friendly manner surfaced. The excellent results obtained from the Project Sydney prototype and the success that the airlines had had with computerized reservations systems suggested an approach. With this in mind, Brandon, by now senior vice-president for planning and information systems (including electronic data processing), sought out and hired Howard Bedford, an expert in airline reservation systems, to develop a comprehensive system. Bedford hired James Tollefson, who had experience at IBM and Avis; Henry Howell; and a crew that soon reached about 150 people, most of whom were also pioneers in high-speed, on-line transaction processing systems. The group was ensconced in the Advanced Systems Development Center (ASD) located in Colorado Springs. The result was COSMOS.

COSMOS (Customer, Operations, Service, Management Operating System) is a mainframe-based order and dispatching system that FedEx has constantly upgraded to cope with increasing volumes, an expanded route structure, and an ever increasing emphasis on tracking packages from origin to final destination. Version 1

began operations early in 1979. It was designed with the idea that the system eventually would be used to keep track of every single package throughout its life in the system. The source of this vision, of course, was Smith.

From the very beginning, Fred Smith wanted to maintain constant surveillance on every package that entered the FedEx system. "Absolutely, Positively, Overnight" was his business battle cry and still is. Thus the continual tracking of packages as they move through the FedEx system is a fundamental, critical success factor of the business.

Early in 1974, Smith had come across a gas station machine that convinced him that a package tracking system might be technologically feasible. FedEx designed a system called "Star" using a gas station imprinter and multi-part forms. Unfortunately an experiment run in Chicago was an abject failure, largely because of the amount of paper and paper shuffling required to handle each package. The Star experiment was squelched; but the dream of 100-percent tracking lived on.

By 1980, with COSMOS and the SuperHub system in place, Smith renewed his demands for tracking, and ASD went to work on it. The result was COSMOS IIA—the first truly positive tracking and status information system—which went on stream in February 1981.

Version IIB—a real-time, point-of-sale to point-of-delivery system with distributed process control—followed in January 1986. Today, with COSMOS IIB, a courier uses a handheld computer (the SuperTracker) with a wand to scan a bar-coded label when a package is picked up, and the vital

information is uploaded to COSMOS. This scanning is repeated five or more times at various stages of handling until the package is finally delivered. The last wave of the wand records the delivery to its intended recipient. A companion system called DADS (Digitally Assisted Dispatch System) receives shipping order information from COSMOS, sends it to a city's dispatch center computer, and routes it to the appropriate courier. The request is displayed on a small DADS video screen in the courier's van or on a portable unit the size of a small briefcase. COSMOS, and to some extent DADS, are predicated on the results that came out of early OR studies.

Releasing the "Butterfly"

Fred Smith characterizes his company as having gone through three entomological phases. About four years beginning in 1973 mark its caterpillar phase during which its real potential and future were fomenting within itself. A short 11 months beginning in January 1977 constitute its chrysalis period, when Smith—with the aid of data obtained from operations research studies—struggled with legislative committees and government agencies in Washington to rid the company of its regulatory constraints. As with any lepidopteran, there was a great deal of activity going on inside the cocoon—in this case considerable OR and marketing planning, but little of it could be seen from the outside. Finally, there is the current butterfly era, which began in November 1977 when the Air Cargo Reform Bill was signed into law, and FedEx was set free.

The butterfly period really began to take off early in 1979, and it presented many new challenges and opportunities for oper-

ations research. The existing package-handling system, young as it was, was already coming to the end of its useful life. Nightly package counts exceeded 50,000 and were growing rapidly. How long would the SuperHub concept remain optimal? What kinds of air cargo carrying capacity would FedEx need in the future? How fast must planes fly to meet their tight deadlines?

Joe Hinson and his group set out to answer these questions. Using a multi-period model of the entire FedEx system and looking five years out into the future, they first estimated the system's carrying capacity, the maximum amount of freight that the system would have to carry each night. One million pounds of total capacity appeared adequate and even provided a little breathing room. (This seemingly conservative maximum load factor was exceeded earlier than anticipated.) They developed several scenarios by making different assumptions about package counts, their distribution, transportation modes (types of aircraft and trucks), and alternative hub-and-spoke structures. They calculated costs and performance factors for every activity: the original shipper request, package pickup, local transport, outbound station sort, outbound line haul (air or ground), hub sort, inbound line haul, inbound station sort, local transport, and final delivery. Once every few weeks, Hinson would trek across to headquarters and present his team's preliminary results to Smith and the executive committee. Some of the executives—almost always Smith—would think of some new possibilities. "What if this were the case?" or "Couldn't we do that?" Hinson dutifully took notes, and as each session came to a close, he would return to

his team to begin a new set of analyses.

This give-and-take dialogue lasted about six months during late 1979. Slowly, a few crucial yet robust results emerged. One was that the company could save on fuel expenditures if it employed DC-10s on high-volume routes to such cities as Boston, Chicago, Los Angeles, and Newark. At first, Smith balked at this solution. For once the company was financially sound enough to make the investment, but he thought the price was too high. One DC-10 cost nearly

Everything is tied to everything else.

as much as his entire fleet of Falcons had cost. The OR team changed its approach. Using an updated version of Autoroute, the team calculated the price point at which FedEx should be indifferent between purchasing DC-10s and another bundle of aircraft. The results rekindled Smith's interest, and, armed with this new information, he made an offer for four DC-10-10CFs. The seller initially demurred. But, because he understood the economics of the deal very well, Smith held tight throughout the negotiations and finally got a purchase price the model told him he could live with. In March 1980, FedEx christened its first DC-10.

Another result of the butterfly explorations was even more controversial, because it reopened old wounds from a previous internal company battle. Various model runs predicted that sometime during 1987 the SuperHub concept would finally reach its practical limit. FedEx would need more hubs. As long as the 727s could handle the

volume, the single hub design remained superior. The DC-10s extended its life somewhat but eventually, as volumes expanded, the SuperHub would collapse of its own weight.

Multi-hubs were in FedEx's future, but when, where, and how should they be operated? The models showed that the most efficient approach was somewhat different from that proposed during the first corporate debate. The original serial approach called for collecting all of a region's packages in a regional hub, fully sorting them there, and then sending packages to destinations within the region while routing all other packages to the appropriate other hubs. This placed a heavy sorting burden on the system. A more efficient approach, it turned out, was for each pickup station to make a simple binary decision. Sending packages for redistribution within the region to the regional hub and sending all others to the Memphis SuperHub. This approach—called the overlay hub method—distributed the sorting burden and permitted it to take place in parallel rather than serially.

The team then used economic location models to determine the best sites for the overlay hubs. Early results showed that the West Coast was a good candidate for an overlay hub, the center of gravity resting just north of Los Angeles. Autoroute runs, however, modified this conclusion. Because of the flight times from Seattle, Portland, and other Pacific Northwest cities and a few other operating considerations, Autoroute indicated that the San Francisco Bay area was a better spot. In June 1986, FedEx opened the first overlay hub in Oakland, California. This was at least a year earlier

than the models had indicated would be optimal. However, Smith was convinced. He wanted to move ahead quickly and to test the concept for real. It has been successful.

Locating the Indianapolis overlay hub was similar. The location model favored Rockford, Illinois. This assumed that FedEx would have to build its own facilities for each of its overlay hubs as it had done in Oakland. In August 1987, however, a fully furnished airport facility became available in Indianapolis. A few quick runs of the models showed that the Indianapolis site was economically viable. Smith bought the Indianapolis facility and began operations in October 1988.

The transportation models uncovered another advantage of the overlay hub concept. The economics of hauling show that for short distances trucks are the most efficient carrier. As the distance increases a switch-over point is reached at which air cargo becomes less costly. The original distribution and configuration of stations favored air cargo. Nevertheless, Smith and his associates had been looking forward to using trucks in high volume areas. Autoroute showed that FedEx had reached that time; the density of traffic around the new overlay hubs made trucks more economical for certain short hauls. FedEx initiated truck service to replace other modes on these routes. Today FedEx is one of the nation's largest integrated, multi-conveyance freight carriers.

Some Current Projects

After extensive studies by the OR team in July 1991, FedEx ordered 25 A300-600 Airbus freighters. The first deliveries were in April 1994. An updated version of Auto-

route shows that the faster cycle time, higher payload, longer range, and capacity for intensive operations should make this aircraft effective on high volume and international routes.

Meanwhile, the OR team is focusing on the least-studied part of the total system—ground pickup and delivery in the field. This activity accounts for about 55 to 65 percent of the total cost of package-handling services.

An experimental courier route planning system project is underway. FedEx is using geo-coding and geo-positioning technology to plot a courier's pickups and deliveries on a map route by longitude and latitude. One use of the map is to show the courier the exact route she followed and to suggest possibilities for more efficient sequencing. Soon to follow are shortest time path estimates. These will help station managers to plan the work of couriers under their supervision. Preliminary estimates show that a five to 10 percent productivity improvement is possible, and almost all of it goes to the bottom line.

Another application involves plotting an entire station's deliveries by route number on the station's service area map. This will reveal operating problems in the route structure. Since customer volumes are always changing, and new customers appear and old ones disappear, route balancing is a constant challenge. Turfs, as couriers' routes are called, are currently plotted by hand, but when the geo-coding experiment is completed, it is likely that a computer program can be devised to do the task more effectively. Ultimately, the map information may be integrated into DADS.

Lessons to Be Learned

The Federal Express Corporation has succeeded by applying the scientific method to its operations. Models and analysis have informed many of its crucial, business-shaping decisions. When the company didn't use OR, beginning with its very first trial run, it performed poorly. Why is this? What explains FedEx's success? What lessons can we learn?

OR has been so successful at FedEx for five major reasons:

(1) The airline and air cargo industries are whole systems businesses with high levels of interdependencies among their parts. They are not easily broken into little pieces to be managed separately. In effect, everything is tied to everything else. These interdependencies are exacerbated by the demands of "Absolutely, Positively, Overnight" delivery. Because of the underlying systemic structure of the business, its route structures and operations are amenable to modeling and analysis.

(2) Fred Smith, the founder and CEO, is devoted to the use of scientific inquiry in his business. In the beginning, he sought out Charles Brandon, listened to him, and gave him responsibility. Brandon became the maestro who translated Smith's vision into OR studies and information and communications systems. With Smith's approval, Brandon hired and initially directed Ponder, Hinson, and others who eventually became part of a world-class team. Ultimately, several senior corporate vice-presidents came up through OR.

Throughout the company's history, Smith has supported the OR effort. The acid test came early on. In 1973, the company had calamitous cash-flow problems. It

asked employees not to cash their paychecks and put all vendors on hold.

National CSS Time Share threatened to cancel FedEx's computer time unless it made an \$8,000 payment. Since the OR effort required the computers, Smith came up with the money. Of equal importance, Smith has always engaged the OR team in debate on issues of crucial importance to the firm. OR has a voice at FedEx.

(3) From the beginning, the OR team focused on issues that were crucial to the business, such as equipment investment and route structure. The team always tried to take the simplest, most straightforward modeling approach it could, as long as it could come up with appropriate and actionable answers. It did not get sidetracked in the esoteric or bogged down in minutia.

(4) The models and their results were considered to be living things, not imperial edifices cast in concrete. The OR team's pride resided more in the company than in a particular model. Through these continual dialogues, the OR team and executive management jointly arrived at shared understandings of the problems they faced, the pros and cons of the models, and the strengths and weaknesses of the model results. They made trade-offs and compromises to achieve FedEx's overall corporate goals. The models informed these judgments. Consequently, when Smith and his executives made decisions, they knew why.

(5) Finally, the OR team conducted itself as a learning organization, a center of knowledge and inquiry about the most important dimensions of the business. Captured in the models and databases are millions of items of information—facts and assump-

tions—that characterize all of the important activities of the business: marketing, route composition, package handling, hauling, and finances. The OR team is constantly updating these facts and assumptions as it discovers new information about technology and operations and new executive insights and desires.

Acknowledgments

This paper is a product of the Harvard Business School MIS History Project. The project is led by James McKenney of Harvard. The team comprises Walter Carlson, Richard Canning, George Glaser, and Ted Withington, who have been associated with information systems and operations research for nearly 40 years and are currently consultants to industry. The other academic members are Richard Mason of Southern Methodist University and Duncan Copeland of the University of Western Ontario. Some of the findings of the project are reported in *Waves of Change* [McKenney, Copeland, and Mason 1995].

We based this article on information gleaned from public sources and from interviews with Jim Barksdale, Charles Brandon, Harry Dalton, Carolyn Freeman, Joe Hinson, Dennis Jones, Tom Martin, Jim Moore, Martha Outlaw, Ron Ponder, Fred Smith, Mike Staunton, S. Tucker Taylor, Jim Turpin, Edwin White, and Denise Wood, all of whom work for or formerly worked for Federal Express. Carolyn Freeman also helped schedule interviews and provided editorial suggestions. Charles Brandon, Joe Hinson, Dennis Jones, and Fred Smith provided comments on an earlier draft. Nevertheless, the authors accept responsibility for any errors in the manuscript.

FEDERAL EXPRESS

APPENDIX: FedEx Summary of Key Events

Date	Situation	Information Source	Model	Decision or Use
March 1973	-Original idea	-Founder	-Initial 11-city model	-11-city schedule (failure)
April 1973	-Soul searching and rethinking	-Emplenment statistics -Employment SIC -Business population -Experience	-112-city origin-destination model	-26-city system (satisfactory)
June 1973	-Cost pressures -Draw on the results of the O/D model	-Airport data -Aircraft characteristics -Loading-unloading -Ground operations -Weather experience	-FLY -Simulation model produces expected flight times and engine cycles	-Refine schedules for more efficient operations
Early Fall 1973	-Desire to grow -Need for additional capital -Investor confusion	-Additional data on candidate cities	-FLY and O/D models simulate 82-city system with 33 Falcons	-Successful presentation to potential investors -Add Los Angeles to system early
Mid Fall 1973	-Need to add financial management to operations management	-Accounting data -Financials -Cost data	-Financial planning model (Fin)	-Refine operations to improve financial performance
Late Fall 1973	-OPEC rationing imposed	-Fuel consumption and cost estimates	-FLY, O/D, Fin models predict five-year fuel requirements for an 82-city system	-Obtain fuel allowance during OPEC restrictions
Summer 1974	-Need for engine overhauls threatens to crnple fleet	-Flight time data -Number of flights -Engine cycles -Overhaul requirements	-Joe/Engine model	-Kept fleet in service until bluckets problem was solved
Summer 1975	-Company growing faster than ability to manage well	-Expected growth -Seasonal factors -Actual load data -ATA gross volume	-Forecasting model, FLY, and other models to predict operating requirements	-Capacity alert bypass hub established at Pittsburgh
Winter 1975	-Increased complexity of the system	-Load factors -Cities served -Arrival and departure times	-AUTOROUTE	-Air Cargo Reform Bill, -FedEx orders seven 727s
December 1975	-Customer calls overload system	-Customer call data -Telephone volume -Telephone facilities -Operating data	-Initial telephone answering model	-Need for prototype centralized system established -Project Sydney

March–April 1997

MASON, MCKENNEY, CARLSON, COPELAND

Early 1978	-Prototype centralized system		-Customer service model	-Decision to initiate COMOS and centralize call centers
Early 1979				-COSMOS Version I implemented
Spring 1979	-More operating problems	-System structure -Cost of aircraft -Facilities characteristics	-Expanded Autoroute model	-SuperHub decision
March 1980			-Autoroute (and other models)	-Add DC-10's to the fleet
February 1981			-Autoroute (and other models)	-COSMOS Version IIA implemented
Winter 1982			-Autoroute (and other models)	-10:30 AM guaranteed delivery decision
January 1986			-Autoroute (and other models)	-COSMOS Version IIB implemented
June 1986			-Autoroute (and other models)	-Overlay Hub established in Oakland, CA
Spring 1991			-Autoroute (and other models)	-Additional overlay hubs -Trucking justified
April 1994			-Autoroute (and other models)	-25 A300-600 airbuses ordered

References

- Sigafoos, Robert A. and Easson, Roger R. 1988, *Absolutely, Positively, Overnight! The Unofficial Corporate History of Federal Express*, St. Lukes Press, Memphis, Tennessee.
- McKenney, James L.; Copeland, Duncan C.; and Mason, Richard O. 1995, *Waves of Change: Business Evolution Through Information Technology*, Harvard Business School Press, Boston, Massachusetts.