

50th Anniversary Article

Improving Emergency Responsiveness with
Management Science

Linda V. Green, Peter J. Kolesar

Graduate School of Business, Columbia University, New York, New York 10027
{lvgl@columbia.edu, pjkl@columbia.edu}

While the goal of OR/MS is to aid decision makers, implementation of published models occurs less frequently than one might hope. However, one area that has been significantly impacted by management science is emergency response systems. Dozens of papers on emergency service management appeared in the OR/MS literature in the 1970s alone, many of which were published in *Management Science*. Three of these papers won major prizes. More importantly, many of these papers led to the implementation of substantially new policies and practices, particularly in policing and firefighting. Much of this work originated in New York City, though many other cities subsequently adopted the resulting models and strategies. In this paper, we look at the context, content, and nature of the research and the factors that led to these early implementation successes. We then track the extent to which these original models are still affecting decision making in emergency response systems. We also examine the pace of development of new OR/MS models and applications in the area. Finally, we look at issues in emergency responsiveness that have emerged recently as a result of the national focus on terrorism and discuss the potential for future OR/MS modeling and application.

Key words: applications; emergency services; fire; police; public sector; urban

Introduction

In his editorial mission statement for this journal, Wally Hopp stated that “*Management Science* needs to play a leadership role in applying our legacy of powerful analytic tools to high-level, long-term planning issues faced by managers” (Hopp 2003, p. v). One area that has enjoyed considerable success in this regard is emergency response systems. Beginning in the late 1960s, papers on the allocation and deployment of police, fire, and ambulance resources that provided important insights, policies, and procedures for managers began to appear in this journal with regularity. This activity continued through the 1970s and into the 1980s. A large fraction of the models in these papers were actually implemented, particularly in New York City, which sponsored much of the basic research; and, many of the models and resulting policies were subsequently used in other cities and had lasting impact on practice.

As the public sector applications department of *Management Science* is the only department of the journal that contains the word “applications” in its title, we thought it fitting that this paper, written for the journal’s 50th anniversary, focus on this era of application success, examine the factors that led to it, and trace the legacy of these publications and models on practice as well as on research. Our intention is not to provide a comprehensive literature survey (see

Kolesar and Swersey 1986, Swersey 1994) but rather to give a more personal overview of the developments and impact of this body of work, taking advantage of our own involvement in its history. We also highlight the role of *Management Science* itself in publishing much of the best of these analyses. In doing so, we hope to provide some insight on the elements of successful model development, implementation, and dissemination in the public sector. And, we feel that it is equally important to examine how and why some of these once-successful models have faded from use while others continue to be implemented, although sometimes in limited ways. Finally, we observe the relative scarcity of papers on emergency services in recent years and ask whether the new attention on homeland security and emergency preparedness may presage renewed interest and activity in model development and application in this area.

History

The late 1960s was a time of unrest in the United States. National crime statistics were rising steadily and Barry Goldwater, the Republican presidential candidate in the 1964 presidential election, made “crime in the streets” a major campaign issue. The assassination of Martin Luther King in 1968 and the civil unrest that followed, the takeovers of a

number of universities, including notably our own home institution, Columbia, and the sometimes violent and always turbulent protests over the war in Vietnam, impressively exemplified in the massive protests at the Democratic Party's national convention in Chicago 1968, were all features of the times.

Simultaneously, in some circles, there was optimism about the potential of computer models and mathematical analysis for solving public policy issues. The Apollo space missions were underway, culminating in the first manned lunar landing, in 1969. Many commentators asked in full earnestness, "If we can land a man on the moon, why can't we...?" The implication was that similar analytical thinking and technology should be used to attack fundamental social problems. The field of OR/MS was still young and many of its practitioners and leaders had either participated in the original World War II military work that the field was rooted in, or had been trained by those pioneers. They shared the heady feeling that mathematical modeling could be applied to many of the nation's problems.

In the domain of public safety and emergency services, two political initiatives were particularly influential and productive. The first we mention only briefly. In response to Candidate Goldwater's charges about crime in America, President Lyndon Johnson in 1965 established The President's Commission on Law Enforcement and the Administration of Justice. The Commission's Science and Technology Task Force was chaired by Alfred Blumstein, then of the Institute for Defense Analyses and an influential figure in the OR/MS community, having served as president of ORSA, TIMS, and INFORMS. The Science and Technology Task Force was charged with exploring how computer modeling and technology could be utilized in attacking crime. Only a small part of its work was related to operational policing issues, but that little amount was noteworthy in that it jumpstarted the career of Richard Larson, then a graduate student at MIT, whose work with the commission contributed to his Lanchester Prize-winning book *Urban Police Patrol Analysis* (Larson 1972). (Larson has also served as president of ORSA and INFORMS.) Overviews of the work of the Technology Task Force on the modeling of criminal careers, the impact of incarceration, etc. may be found in the papers of Blumstein (2002), Larson (2002), and Maltz (1994).

The second political initiative, The New York City–RAND Institute (NYCRI), had a more profound impact on the development of emergency service deployment modeling. The institute was a unique partnership between the City of New York and the RAND Corporation of Santa Monica, California. RAND was the original and prototypical military think tank; a direct outgrowth of the first operations

research work done for the U.S. Air Force during World War II.¹ By the late 1960s, RAND had started to broaden its portfolio of work into domestic policy research in areas such as telecommunications, education and human resources, and energy.

Here is a thumbnail sketch of how and why the NYCRI came about.² John V. Lindsay was elected mayor of New York City in 1966 on a platform of governmental reform. To this end, his administration carried out substantial structural city government reforms, including a major reorganization that reduced the number of agencies reporting directly to the mayor from 50 to 12. In what was at that time an original move, Lindsay hired a core of some 50 "analysts/administrators." E. S. (Steve) Savas, originally at IBM and the author of several *Management Science* papers on public sector issues, was one of these analytically trained administrators (Savas 1969, 1973, 1978).

Fredrick O. Hayes, Mayor Lindsay's first budget director and a key figure behind these reforms, was motivated to bring even more analytic competence to bear on city problems with his perception of the "appalling growth rate of virtually all of the problems to which municipal programs and services were directed. Crime, drug use, fire alarms, solid waste..." (Hayes 1972, p. 1). This dismal prospect and his faith in analysis lead Hayes to request Ford Foundation funding to create a RAND-type institution devoted to the city's problems. When this did not materialize, he went to RAND itself. This chance to diversify out of military work was exactly what RAND was looking for and negotiations led to the formation of the RAND Corporation–New York City partnership in late 1968. While NYCRI would eventually work on diverse problems in many city agencies, the initial research contracts were with the Fire Department, the Housing and Development Administration (largely public housing and administration of the city's rent control laws) the Police Department, and Health Services (primarily the administration of the municipal hospital system).

While the institute's formal status with the city was that of a contractor, from the outset both the city and RAND saw this as more than the usual consultancy. The relationship with RAND was to be long term and comprehensive, focusing on immediate problems as well as on those that would require sustained study, experimentation, change and reevaluation. Mayor Lindsay spoke of the "...willingness

¹ For one of the few examples of RAND's military work that is unclassified, see Wohlstetter et al. (1954).

² More details can be found in Drake et al. (1972), particularly in Chapters 1, 2, 7, and 9; in Greenberger et al. (1976), particularly in Chapters 7, 8, and 9; and in Walker et al. (1979, pp. 629–639).

on the part of RAND to leave model building long enough to assist in the application of their new systems in a real agency in a real city" (Dickson 1971, p. 249). On the other hand, NYCRI's first president wrote that the researchers should be "insulated enough from city hall's daily operational concerns to work persistently on underlying problems" (Szanton 1972, p. 20).

However, the political landscape of New York was not simple, and having the mayor and his budget director as enthusiastic clients did not itself make the going easy. By hiring RAND without their prior consultation, the mayor had in effect imposed RAND on the city agencies. Some in the agencies saw the RAND researchers as spies for City Hall. One of RAND's political scientists and the second leader of the RAND Fire Project observed, "Since the consultant must have information . . . and often only the bureaucracy has that information, by selectively cooperating with the consultant, the bureaucracy can effectively scuttle the desired change" (Archibald and Hoffman 1969, p. 10). The city's chief financial officer, Comptroller Abraham Beame, had been the mayoral candidate whom Lindsay had defeated. Beame's opposition to the institute was not concealed. He was a supporter of the political status quo and his view, shared by many, was that anything RAND was doing, if it were indeed worth doing, could be done more cheaply by professors at the City University of New York. The City Council, the legislative branch of city government, had to approve the RAND contracts and soon demanded a voice in shaping the research agenda. The city's powerful municipal worker labor unions quickly accused the institute of being little more than stop-watch carrying, antiworker efficiency experts. Before long, left-leaning commentators portrayed the RAND researchers as a collection of cold-blooded Dr. Strangeloves who would have the city burn down at the altar of cost-effectiveness modeling (Hoos 1972, Wallace and Wallace 1999). Indeed, some of this criticism made its way to the pages of this very journal (Wallace and Wallace 1980, Chaiken et al. 1980).

To a considerable extent the RAND strategy was to stay out of the limelight and try to make other people look good. But this could backfire too. After several years of generally successful work, the institute recognized that it had no constituency outside of the mayor's office. Moreover, a credit/blame game developed in which even the successful Fire Project researchers felt that the New York Fire Department (FDNY) chose to take credit for itself, or let RAND take blame according to its own convenience. In this environment, it is remarkable that anything positive was achieved by "systems analysis and management science."

The RAND Fire Project

One of the initial NYCRI contracts with the city was with the FDNY, a project that began at the very start of the RAND–New York City relationship in January 1968 and lasted until the institute's doors closed in September 1975. The Fire Project was, by many accounts, the most successful of the NYCRI ventures. The FDNY's problems were very painful. In the five years from 1963 to 1968, fire alarms in New York City increased 96% from 116,000 to 227,000, while firefighting resources stayed almost constant. FDNY operating expenses were increasing at over 20% per year, largely as a consequence of wage increases for its 14,000 uniformed firefighters. Workloads on individual firemen were excessive, with some fire companies responding to alarms more than 8,000 times a year, or nearly once an hour, 24 hours a day, 365 days a year. During peak times, some fire companies ran from one incident to another all night long. Dennis Smith's best-selling book, *Report from Engine Company 82*, gives a compelling picture of the stress and danger the firefighters faced at this time (Smith 1969).

The fire communications system, still the same telegraphic-driven bell system that had been designed when fire engines were horse-drawn, was becoming severely congested. Everyone's "pet solution" was to bring technology to the rescue: "Let's get a new high-tech communications system for FDNY." This was to be RAND's first mission. But before too long, the researchers found that communications problems, while real, were not quite what they were originally thought to be, nor would an efficient solution be automatically achieved by "a big computer in the sky" (Greenberger et al. 1976, Chapters 7 and 8). While important queuing analyses were being done on immediate communications bottlenecks by project member Arthur Swersey, the RAND team got FDNY support for its view that resource deployment and fire incidence forecasting were fundamental issues that needed intensive study. The Fire Project was initially staffed primarily by RAND military systems analysts from Santa Monica. When it was realized that OR/MS work would be at the core of the project, RAND began to recruit additional management scientists locally. First Ed Ignall and Art Swersey, and then Peter Kolesar and Kenneth Rider, all of Columbia University, joined the project team. The project staff also included Warren Walker, Jan Chaiken, and Edward Blum, RAND-based researchers who were hired directly into the New York City office. Most of the academics would at some point take extended leaves of absence from their home institutions to work on the Fire Project full time.

The nature of the fire problems, the state of the art of OR/MS research on emergency services, and the proclivities of the researchers led to a particular

style of working. First, there was essentially no existing OR/MS literature, so the researchers had to start afresh. (Exceptions, like the papers of Valinsky 1955 and Hogg 1968, were of limited use.) The team, which had a pragmatic orientation, went to fires with the firefighters, slept over in firehouses, sat at the shoulders of the dispatchers in the communications center, poured over fatal fire reports in the archives, and, most importantly, cleaned up the fire data tapes. Although the FDNY had created computerized records of all fire alarms for the five years prior to the start of the project, the data had never been analyzed; the piles of punch cards were in terrible shape and it took months of work to get these records cleaned up and into a format in which they could be used for analysis.

The first major piece of management science work done by the RAND team was the creation of a simulation model of firefighting operations (Carter and Ignall 1970; Walker et al. 1979, Chapter 13). Development of this model received highest priority as real-life experimentation with the fire system, such as changes in dispatch strategies, the communication system, or the number of fire companies on duty, would be too expensive or dangerous. Technical aspects of the design of the fire simulation model were innovative and influential (Carter and Ignall 1975), and it proved to be the foundation on which both specific new deployment tactics (Kolesar and Walker 1974, Ignall et al. 1982) and general new theories (Kolesar and Blum 1973) were validated.

Although the research team eschewed theoretical formulations for their own sake, much of the early work on these issues had a distinctly theoretical OR/MS flavor: queueing models of fire company availability (Chaiken and Ignall 1972, Carter et al. 1972); an empirical Bayes' approach to alarm forecasting (Carter and Rolph 1973, 1974); a stochastically based integer linear programming formulation of fire company relocations (Kolesar and Walker 1974); and Markovian decision models of initial dispatch to a new alarm (Swersey 1982, Ignall et al. 1982). While each piece of work contributed to the researchers' understanding of fire deployment, the sophisticated models were generally not the most important contributions to actual operations. There were two dominant reasons for this. First, many of the problems that intrigued the researchers from an analytical perspective were tactical and had limited impact. One such example was identification of the fire companies that would constitute an optimal initial dispatch to a new alarm of unknown severity. The elegant stochastic formulation that was developed for this rather micro problem only made sense in the most stressed neighborhoods of New York City, as it focused on the phenomenon of random fire company unavailability

that was at the heart of the dispatch dilemma in these neighborhoods. Second, the insights obtained from the more sophisticated models could often be translated into simple heuristics or rules. For example, ideas emanating from the research on initial dispatching were implemented as simple ranked lists of high-priority alarm boxes in the most stressed neighborhoods (Ignall et al. 1975).

Such tactical issues were not as vital to senior FDNY management as the broader questions of how many fire companies were really needed in New York City, and where should they be located to provide fair and adequate protection to neighborhoods as diverse as the Wall Street financial district, the Upper East Side, Harlem, and the suburbs of eastern Queens. These macro issues involving millions of dollars eventually succumbed to much simpler, often deterministic, models. A prime example is the "firehouse-siting model," which was the simplest of all the models developed, but also the most influential and enduring. It took the RAND team a while to understand that simple models would work, though. One could consider the contribution of much of the more elegant work as being partly foundational and partly a side payment to the research team taken in the form of publications in journals such as *Management Science*. It would have been very difficult to keep this talented team together without the prospect of publication in refereed journals as none of the academic members of the team was yet tenured. The research results of the Fire Project were formally reported to the city as a series of RAND Research Reports, or "Rs," many of which are still available in RAND deposit libraries around the world and from RAND's website at <http://www.rand.org>.³

Sometimes theoretical research yielded big payoffs, however. A prime example was the work on fire company travel distances and travel times, which resulted in the square root law (Kolesar and Blum 1973) described in the next section. It was largely motivated by B. O. Koopman's derivation of his logarithmic laws of search effectiveness, which led to a breakthrough in deployment of antisubmarine search forces during World War II. Koopman, a mathematical physicist, had derived his laws of search from physical principles, and they were later confirmed by direct observation (Koopman 1956a, b; 1957). One of the authors, Kolesar, a former student of Koopman's at Columbia, was troubled by the fact that the Fire Project team often resorted to ad hoc detailed calculations of the consequences of alternative policies. After spending countless days in a room at RAND that was wallpapered from floor to ceiling with enormous

³ Searching on the key words "fire" and "police" will lead to most of the publications related to emergency service deployment.

maps of firehouses and fire locations, he started to wonder: “Why isn’t there a simple general law of fire protection that would answer the big questions of where companies should be located and how many there should be?” When Kolesar articulated this musing to Ed Blum, the Fire Project leader, Blum replied at once: “Oh, I know what it is. Response distance goes down with the square root of the number of fire companies.” “How do you know that?” Blum responded, “It’s simple: Distance is the square root of area.” It took months of work to fully develop and test this theory, but in the end it became the cornerstone of many of the most useful analyses done by the team, including its work on the biggest issue of all—determining the number of companies that should be located in each of the FDNY’s commands. A move toward analytic simplicity had begun, but was only credible because of the work that had come before, particularly, the development of, and experiments with, the simulation model.

As response *time*, not response *distance*, was the key proxy used to evaluate the FDNY’s performance, it was important to model fire engine travel velocity. No one inside the FDNY knew how fast engines went and how speed varied by time of day or weather, etc., so the team designed experiments to measure velocities across the city. These were resisted at first by the firefighters’ union—some stop watches were thrown against fire house walls or mysteriously dropped off the fire trucks. But, in the end, the data was collected and provided the basis of a nonlinear regression model of fire engine travel time as a function of distance that has since been revalidated in other cities (Kolesar et al. 1975a).

Some of the practical impacts of the Fire Project on New York City are documented in the project’s application for the 1974 College on the Practice of Management Science (Edelman) Award, including demonstrated annual savings of \$5 million on a base of a \$375 million operating budget. The Fire Project was costing the city about \$500,000 a year—less than a single fire company. As detailed therein, the team’s research played a role in the city’s decisions to close six fire companies and permanently relocate seven others, all carried out in 1972 when alarm rates had fallen off a bit and budgetary pressures were increasing. In addition, an “adaptive response” initial dispatch policy had alleviated workload in the high fire-incidence regions of the Bronx and Brooklyn, and the FDNY was planning to implement a dynamic relocation algorithm. Indeed, by 1975, this real-time, integer LP-based relocation algorithm designed by Kolesar and Walker was successfully implemented as part of a new computerized FDNY Management Information and Control System. During the years when New York City’s alarm rates were at their peak,

the algorithm was used to suggest fire company relocations many times daily.

Over the next several years the pace of work on macro resource-allocation issues quickened as the FDNY was repeatedly called on to contribute cost reductions to the city’s efforts to balance its budget. By 1974, the city had entered a severe budget crisis and more fire company closings were explicitly requested by the mayor’s office. The RAND models were used to identify those company changes that potentially had the least deleterious impact on fire protection. This happened in several waves and, in total, by 1978 24 fire company locations were closed and 10 of the companies permanently disbanded. These closings were challenged in court by the firefighter’s union and the affected neighborhoods. The strong analytic basis underlying the closing decisions was a factor in defeating these lawsuits. At first the analyses were rather ad hoc and employed the square root model to estimate the impact of average response times in the impacted regions. Later, the RAND firehouse-siting model, which was developed in 1975 as the NYCRI was closing, played a central role in the closing decisions (Walker et al. 1979, Chapter 9; Dormont et al. 1975; Walker 1975). Based on a computerized map of all fire alarm boxes and fire company locations in a region, this model computes static response times for each alarm box in that region using historical alarm frequencies and the square root law and travel time models. While many other considerations, some overtly political, came into play, the models were influential.⁴ Moreover, to support the continued use of the RAND models and the analysis mission in general, in 1974 the FDNY created a Division of Planning and Operations Research staffed with OR/MS professionals.

There were frustrations as well. For example, the RAND analysts realized early on that there was a gross mismatch between the constant number of fire companies on duty around the clock and the enormous peaking of both false alarms and real fires in the evening hours. A variety of staffing alternatives were devised by the team to improve the hourly match-up between the number of fire companies on duty and the alarm rate, but the strong opposition of the politically powerful firefighters union defeated them all (Ignall et al. 1975, Kolesar and Rider 1981).

On the management science front, the RAND team produced some 15 papers in refereed journals that were directly based on the project’s deployment research. Of these, four appeared in *Management Science* while five were published in sister journal *Operations Research*. The work of the team garnered

⁴ A contemporaneous viewpoint on this work by the then fire commissioner can be found in O’Hagan (1973).

three awards: the 1975 Lanchester Prize of ORSA for the relocation algorithm (Kolesar and Walker 1974); second prize in the 1974 Edelman competition on the practice of management science (Ignall et al. 1975); and the 1976 NATO Systems Science Prize for a collection of 13 technical publications produced by the team between 1970 and 1975. It is noteworthy that, because of their operational and applied nature and the modest number of new theorems they contained, much of this work was difficult to get published in top-tier journals. For example, the Lanchester prize-winning paper on the relocation algorithm was rejected by *Management Science* because, as noted by one referee, “it contains not a single lemma or theorem.”

After the initial successes of the Fire Project, the NYCRI was able to obtain modest financial support from the U.S. Department of Housing and Urban Development (HUD) for the dissemination of the institute’s research on emergency service deployment to other cities. The methodological legacy of the project, in addition to the articles scattered through the technical literature, is the book produced by the team members under support from HUD (Walker et al. 1979). In this book, the authors linked the various stand-alone journal articles and RAND reports into an organized course on fire deployment analysis. Material was added about the background of fire services, change management issues, and the like. The level of technical discussion of the models was targeted to future fire department staff analysts. However, nothing was oversimplified. It is also noteworthy that the firehouse-siting model, which would ultimately prove to be so useful in New York City, was funded by HUD and was initially applied in cities such as Yonkers, New York; Trenton, New Jersey; Denver, Colorado; and Wilmington, Delaware.

The Police Project

A parallel NYCRI contract was initiated with the New York Police Department (NYPD) when the institute opened in 1968. The Police Project’s original scope was not as focused on operational deployment issues as was the work with the FDNY. Much of the police work in the early years was on “softer” policy analysis, e.g., minority recruitment, effectiveness of criminal investigations, and police corruption. Alone among the Police Project members, Richard Larson worked on deployment issues related to the operations of New York’s 911 emergency telephone system dispatching office. His suggestions were almost immediately implemented (Larson 1972, 2002). Overall, however, there was not the intensive cooperation on deployment that characterized RAND’s relationship with the FDNY. Personalities and NYPD politics played a critical role. Whereas the FDNY chief saw the RAND researchers as allies propelling agendas that were his own, the RAND team was unable to gain an

insider position with the most powerful forces in the NYPD.

During the early NYCRI years, the NYPD was an institution under siege (Murphy 1977, Daley 1973). There was massive pressure on the department over the issue of police brutality toward minorities. Then, the devastating police corruption scandals that surfaced in the Knapp Commission (Knapp 1972) investigations forced the resignation of Howard Leary, the first police commissioner in office during the Police Project. His replacement was a man who appeared to think that he personally had all the answers to the department’s problems. The NYPD as a whole was demoralized and defensive. It was not an environment conducive to research. In addition, it did not help that RAND was associated with Mayor Lindsay, who had tried to impose an effective civilian review board to oversee citizen complaints of police violence. Under hostile examination during a city council hearing, the president of the institute testified that the police commissioner had expected the RAND team to be seasoned veterans but complained that key participants were “young MIT graduates.” Further, he stated that, in contrast to what was happening at FDNY, RAND was telling the police things they did not want to hear (Ranzel 1970).

The so-called “fourth platoon” controversy illustrates how the RAND analysts were stymied in their first year of work with the NYPD. The gross mismatch between how patrol cars were scheduled for duty as compared to the hills and valleys of the daily temporal demand for police services had been noted by the Lindsay administration even before RAND came on the scene. The mayor tried to get the NYPD to implement a corrective “fourth platoon”—an overlay tour on the traditional three “ platoons” that worked midnight to 8 a.m., 8 a.m. to 4 p.m., and 4 p.m. to midnight. RAND researchers actually generated one version of a fourth platoon plan, but only the mayor’s office was receptive to it.

After the first unsuccessful year, police work at NYCRI went into dormancy with no further NYPD funding, but was modestly sustained by HUD. Then, in 1974, optimistic that the Fire Project team’s talents could be equally helpful to the NYPD, RAND reassigned researchers Chaiken, Swersey, Kolesar, and Walker to a renewed Police Project. After new team members rode around in patrol cars, sat in the dispatching center, and scrutinized data from 911 tapes, a detailed simulation model of police patrol operations was developed (Kolesar and Walker 1975). However, the only policy it was used to test—cross-sector patrol car dispatching—was never seriously considered by the department.

During this renewal of the Police Project, the issue of the mismatch between the number of cars actually fielded and the demand for emergency police

service was revisited. Using a combination of optimization and queueing theory, the researchers generated a range of staffing options that were much more flexible than the dormant fourth platoon concept (Kolesar et al. 1975b). However, the concepts were never seriously considered for adoption by NYPD senior management and were formally abandoned after a successful court challenge by the police union (Moore et al. 1975). But the extent of the staffing mismatch was documented as never before, and a fourth platoon program was reinstituted in a scaled-down version staffed by police officers on a voluntary basis.

The most successful of the institute's police deployment models was the Patrol Car Allocation Model, or PCAM (Chaiken and Dormont 1978a, b). This queueing-based optimization model, described in more detail in the next section, created an efficient method for determining the allocation of patrol cars and officers across the seventy-odd police commands (precincts) of the city and across the three tours of duty. Like the firehouse-siting model, it played a critical role in determining the best way to reduce resources during the financial crisis of the 1970s. And, as with the firehouse-siting model, it was largely developed and disseminated to other cities with support from HUD, and was widely distributed to, and used by, other cities.

Though it was never implemented as part of the Police Project, the hypercube model (Larson 2001) was also partially funded by the NYCRI. It was used later in a study of travel times for the NYPD (Larson and Rich 1987), as well as to support the deployment of ambulances and police cars in various cities including New York, Boston, and Orlando (Brandeau and Larson 1986, Sacks and Grief 1994).

Management Science 1969–1989: A Bounty of Applications to Emergency Systems

Management Science published some of the earliest and most practically influential papers in the area of emergency response systems, many of which were a direct or indirect product of the New York City-based work described earlier. Here, we focus on these as well as other papers that described or resulted in implementations elsewhere.

Perhaps the earliest of these is a paper by Savas (1969) on a simulation analysis of the ambulance system in a single hospital district in Brooklyn exploring the potential improvements from proposed changes in the number and location of ambulances. This marked the first time that New York City used simulation as an aid in decision making. Before Savas' study, ambulances were located at each district's hospital. To reduce response times, a proposal had

been made to station ambulances at satellite garages located in the middle of the highest demand areas. The simulation study indicated that this would substantially improve response times, and as a result a satellite was placed in the test district on a pilot basis. But a new and more fundamental recommendation emerged as well. It became evident that ambulances should be stationed close to the demand and not tied to hospital locations. This observation—that the transportation service could be divorced from the treatment centers—implied that ambulances should be centrally dispatched and managed, that they should be dispersed throughout the hospital district, and that they should be relocated as demand patterns change. Largely as a result of this work, New York City changed its policy and began to locate ambulances at curb sites, a practice that continues to this day.

This pioneering use of simulation and the resulting practical policy implications were a major impetus to the use of simulation and other quantitative modeling in emergency vehicle location. Among these was another early and influential *Management Science* publication by Fitzsimmons (1973), which focused on identifying optimal ambulance locations. Using an $M/G/\infty$ queueing model combined with simulation, he estimated the probabilities of particular ambulances being busy assuming a given set of possible locations. This methodology was coupled with a pattern search routine in a computerized ambulance deployment model named CALL (Computerized Ambulance Location Logic), which identified ambulance locations that minimized mean response time. CALL was used successfully to choose 14 out of 34 possible firehouses at which to station ambulances in central Los Angeles and resulted in significant improvements in response times. It was also used to plan an emergency ambulance system for Melbourne, Australia. Another application of simulation to locate ambulances was described in an early *Management Science* paper by Swoveland et al. (1973), who developed a simulation model of the ambulance system in Vancouver, Canada, to estimate mean response times and other performance statistics for various possible ambulance locations and dispatch policies. They then used this output in a combinatorial optimization model to identify near-optimal ambulance locations.

To identify emergency vehicle locations that minimize total mean response time, it is necessary to estimate the average travel time as a function of any particular set of available units. It had already been demonstrated (Larson 1972) that the average travel distance in a region is inversely proportional to the square root of the number of available units per unit area. So, by using a queueing model to first obtain the probability distribution of the number of busy

units in a region with N units, the square root model could be used to estimate distance, and hence travel time, for each possible system state. However, in a pivotal *Management Science* paper, Kolesar and Blum (1973) showed that the average travel distance is also approximately inversely proportional to the square root of the *average* number of available units per unit area. Thus the expected travel time could be estimated simply without the need of a queueing model. As mentioned previously, this square root model was used extensively by the RAND Fire Project, particularly to identify which fire companies to close and where to relocate others to improve response times (Kolesar and Walker 1974), as well as in many subsequent papers on emergency response planning and management (see, e.g., Swersey 1982, Ignall et al. 1982, Green and Kolesar 1984b, Halpern 1979).

The Kolesar and Blum square root model was the foundation for another important *Management Science* paper on allocating fire companies. Rider (1976) used the model to enable managers to incorporate non-quantifiable criteria into decision making about fire company allocations. Rider's Parametric Allocation Model used a parameter in the objective function of an optimization routine to represent the trade-off between minimizing citywide average travel time and equalizing average travel times across regions, where travel time was calculated using the Kolesar and Blum formula. This provided a more powerful tool for fire department managers, who could now consider a range of allocations and use personal judgment to choose one to achieve a desired balance of efficiency and equity. The Parametric Allocation Model has been used in several cities including Jersey City, New Jersey and Tacoma, Washington (Walker et al. 1979, pp. 349–364, 581–588).

Ambulance and fire systems were not the only beneficiaries of work that was published in *Management Science* in the 1970s and 1980s. Some of the most influential research was in the area of police patrol. One of the most widely disseminated models in the area of emergency responsiveness is the Patrol Car Allocation Model (PCAM), which was described in two *Management Science* papers by Chaiken and Dormont (1978a, b). PCAM, which was part of the NYCRI work, was developed as a result of senior NYPD management's interest in developing a quantitative, independently justifiable method for allocating police personnel to precincts. The patrol force allocation method generally favored before the RAND work used a subjectively weighted average measure of various disparate factors considered important by police departments in determining staffing levels, including precinct sizes, crime rates, and numbers of arrests. However, Larson had showed that these "hazard model" formulae did not actually work the

way police commanders thought; PCAM's structure was a direct outgrowth of his work (Larson 1972, §1.4 and Chapter 5). Although various queueing-based models had already been developed for patrol allocation in New York City, St. Louis, Los Angeles, and Rotterdam, each had limitations that precluded its general usefulness. PCAM, which was designed after a review of these earlier programs, used an $M/M/c$ queueing model with priority classes and could operate in either descriptive or prescriptive mode. It provided a variety of output measures, including the average queue time by priority class, the average travel time, patrol car utilization, and preventive patrol frequency. Average travel time was calculated using the square root formula of Kolesar and Blum (1973) described above, while the preventive patrol frequency used a formula developed by Larson (1972). It allowed an adjustment for noncall for service work, including activities such as meals, auto repairs, and special assignments, which was found to account for as much as 60% of total patrol time in some cities. In prescriptive mode, PCAM allocated car-hours to shifts where a shift is a combination of a specific tour of duty on a specific day in a specific precinct. This allowed users to implement tours with differing lengths. It also allowed for "overlay" tours, beginning during one standard tour and ending during another. PCAM could either determine the minimum number of cars needed in each shift to meet user-specified performance constraints, or allocate a fixed number of car-hours among precincts for a given shift or among shifts to minimize a given objective function. These features, which allowed police managers to specify inputs and outputs in ways that were meaningful to them, is what made PCAM so valuable and widely adopted.

PCAM was originally validated using data from New York City and was used during the financial crisis of the 1970s to make difficult decisions about cutbacks on patrol resources. It was ultimately distributed to over 40 police departments in the United States, to cities in Canada and the Netherlands, and to the single police force which covers all of Israel. In most of these locales, PCAM was implemented after parts or all of the model were validated using local data, and its use led to operational changes (Chaiken 1978, Lawless 1987).

One significant shortcoming of PCAM was that it did not explicitly represent multiple-car dispatches. Every police department receives some calls that require the services of more than one patrol car. In New York City, for example, over 30% of calls result in a multiple-car response. When the size of the patrol force relative to the call rate is large, a simple upward adjustment to the call arrival rate by the multiple-car dispatch ratio may result in fairly accurate predictions of delays. However, after the size of the

New York City patrol force was reduced in the late 1970s, the NYPD found that despite such heuristic adjustments, PCAM was significantly underestimating actual delays (Green and Kolesar 1984a, 1989). This led them to contact Peter Kolesar to commission a revision of the model. Kolesar, in turn, enlisted his Columbia colleague, Linda Green, who had recently developed a queueing model in which the number of servers needed by a customer was random (Green 1980), and this resulted in Green's development of a multiple-car dispatch (MCD) queueing model of police patrol published in *Management Science* (Green 1984). The MCD model is a multiserver, multipriority Markovian queueing model in which the user specifies a probability distribution of the number of servers needed by each call for service type. In an extension of the basic model, both a minimum and maximum number of servers may be specified and the actual number used is dependent on server availability. In the MCD model, service to a job does not begin until the minimum number of required servers is available. Once service has begun, service times of cars are identically and independently distributed. Various performance measures are computed, including the probability of delay and mean delay by priority class and the average number of available servers. The MCD model was validated in New York City and incorporated in a revised version of PCAM (Green and Kolesar 1989, Chaiken et al. 1985) that was distributed through RAND to 46 police departments. It was also used in the evaluation of the proposed mergers of the police and fire departments in several cities (Chelst 1990) including Grosse Point Park, Michigan (Chelst 1988), where the analysis demonstrated that a merger could bring improvements in response times, patrol coverage, and operating expenses, and convinced voters to support it in a referendum.

The MCD model also played a central role in what would become a politically controversial study commissioned by the New York City mayor's office in 1981 to determine whether the NYPD should switch from two officers per patrol car to one. This study is described in another *Management Science* article (Green and Kolesar 1984b). The city was pursuing a gain-sharing program in negotiations with the police officers' union and wanted to determine how many more one-officer cars should be fielded to achieve the same average dispatch delay as with the current two-officer system. The Green and Kolesar study, employing the MCD model, showed that though a one-officer patrol system could achieve a significant reduction in police officers, about 40% more patrol cars would need to be fielded. This increase of cars in the street and "on-the-air" raised concerns on the researchers' part about the capability of the

existing 911 management and communications system to coordinate the back-up car dispatch needed to assure patrol officer safety. No one, either in the mayor's office or in the NYPD, ever challenged the accuracy of the researchers' findings or conclusions. Indeed, senior NYPD commanders concurred with these concerns, and also with Green and Kolesar's suggestion that a carefully monitored experimental one-officer program be conducted in a limited number of precincts. The mayor's office, however, eager to demonstrate, or at least claim, productivity gains before an upcoming election, decided to pursue the program with neither further study nor a program to address the potential communications bottlenecks. When Green and Kolesar refused to support the city's intention to implement a broad one-officer program, the mayor's office withheld payment on part of their completed work and threatened a lawsuit to obtain the then unpublished MCD model. The mayor's office intended to provide the model to another OR/MS analyst who, it hoped, would support its position before the City Council. However, negotiations broke down after the police union's leaders, aware of the concerns raised in the Green and Kolesar analysis, insisted that the city guarantee that the number of cars fielded be adequate to assure rapid back-up, and the city refused. Over the next several years, the results of the Green and Kolesar study were used to justify the city's further exploration of one-officer patrol, which was ultimately implemented on a limited basis. (Green and Kolesar, however, though independently funded to continue their research on patrol deployment, were blacklisted from working for the city until the next administration came into office.)

In another influential *Management Science* paper, Chelst (1981) compared one- vs. two-officer patrol systems by estimating the differences in travel times for both first- and second-arriving units using an approach similar to Kolesar and Blum (1973). Chelst considered two different models of dispatch operations: a conventional "beat" system in which cars are assigned to a particular geographic region and, when available, respond to all calls originating in that region; and a system in which an automated vehicle monitoring system allows for cross beat dispatch of the unit closest to the incident. An important contribution of this paper was a clarification of workload conditions under which two one-officer cars could arrive at the scene faster than one two-officer car. In early 2003, the city of Buffalo, New York, facing bankruptcy, was considering a switch from a two- to one-officer patrol system with a 20% reduction in officers. Because of his *Management Science* paper, Chelst was asked to assess the impact of this proposed change and develop a patrol deployment plan.

The subsequent study and testimony before the City Council led to a negotiated agreement with the police union to switch to a one-officer system in the summer of 2003 that resulted in a dramatic improvement in patrol response times (Warner 2003).

What's Happened in the Last 15 Years?

In total, between 1969 and 1989, over two dozen articles focusing on emergency response systems appeared in *Management Science*, many of which described new models that influenced actual operating policies and practices. Since that time, we were able to find only two articles in the journal on emergency systems (Rajan and Mannur 1990, Athanassopoulos 1998), neither of which appears to have resulted in an actual application. This is a consistent with a general lack of papers in this area across all of the management science/operations research journals starting in the 1990s. Why did this happen? Does this dearth of publication activity imply that there was a sharp decline in interest in, or need for, models to aid decision making in ambulance, fire, and police systems? And what, if anything, does this imply for the usefulness and use of the numerous models that were developed and implemented in the 1970s and 1980s?

There are several possible explanations for this decline in emergency response system publications. To some extent, the models developed in earlier decades had already addressed the most basic and important problems faced by the managers of emergency systems. In effect, much of the cream had been skimmed. Many of these models were widely known and disseminated through the efforts of RAND and the publication of two major books describing them: *Urban Police Patrol* (Larson 1972) and the *Fire Department Deployment Analysis* (Walker et al. 1979). Some of the models have been modified over the years by consulting organizations and incorporated into proprietary software packages, precluding open literature publication of these modifications. In addition, although there are gaps and flaws in the original models that could be addressed with more sophisticated approaches, increased complexity could be resisted by the managers and planners who use such models. Typical public sector technology transfer and implementation challenges are described in two studies published in *Management Science* (Chaiken 1978, Lawless 1987).

Another possible explanation is a diminished need for decision-support models in emergency response planning and management. As described previously, in the late 1960s to the mid-1970s large U.S. cities were confronting problems of increasing crime, drug abuse, and social and racial unrest. An economic downturn led to more arson-for-profit as distressed

business owners torched their own premises to collect insurance money; "burn, baby, burn!" was a popular chant that reflected the climate of social protest. During this period, demand for emergency services grew and economic pressures constrained the available resources, making it imperative that emergency systems utilize their scarce resources as efficiently and effectively as possible. In contrast, the 1990s were a time of economic prosperity, low unemployment, and decreasing crime and turmoil. Staffing levels in police and fire departments were increased, and hence the need for "optimal" resource allocation became far less pressing.

The previous discussion highlights what we feel is a distinctive feature of research in emergency response systems: It is very difficult, if not impossible, to do without the sponsorship of a client organization. Unlike research in, for example, inventory management or queueing theory, which is often conceived of, and carried out by, individual academics scattered across the university landscape, meaningful models of emergency systems cannot be developed without intimate knowledge of the organization, its operations, and its objectives. When the need for such models is low and there is little political impulse to identify critical social problems, provide and focus funding, encourage cooperative partnerships, and guide implementation, scholars will find it very difficult to identify and/or implement a fruitful stream of research in the area on their own. One example is our work on staffing service systems that face time-varying customer demands (Green et al. 2001), which was originally motivated by our finding that the NYPD's use of the standard Erlang-based approach resulted in understaffing (Green and Kolesar 1989). This research led to the development of a simple heuristic that corrects the understaffing problem, but it was never implemented in the NYPD.

Moreover, it seems that an element in the success of the NYCRI effort was its large scale, temporal continuity, and unique mixture of consulting and research (Hayes 1972). A stable core team of about half a dozen Ph.D.-level OR/MS researchers stuck with a set of problems from conceptualization to implementation over a six-year period. Moreover, the analysts were amply supported by a staff of computer systems people, data analysts, and a management team that ran interference for them with the political establishment.

To gain more insight into the current need and use of management science models in emergency systems, we examined whether and how the models that were developed by the NYCRI are currently used in the NYPD and FDNY. We interviewed managers in the planning organizations of both departments to determine which models are still in use, how they are being used, and the extent to which senior managers see the

models as useful for the problems their organizations are currently facing.

Our discussion with a group from the NYPD's Office of Management Analysis and Planning (OMAP) confirmed that the department still uses the revised version of PCAM that incorporates the MCD model (see the discussion in the previous section), but has changed the way in which it is used. First, and somewhat ironically, although the MCD model resulted from an explicit request by the NYPD to correct PCAM's inability to adequately account for the high level of multiple-car dispatches in New York City, OMAP does not implement this feature. This is despite the fact that the current fraction of 911 calls with multiple-car response is between 30% and 40%. Second, while the PCAM model was run in the 1980s on a monthly basis to adjust patrol-car allocations and staffing for seasonal changes in 911 demand patterns, its current use is mostly limited to providing an objective first step for "equitably" allocating the annual graduating class of police recruits among the city's 71 precincts according to several performance and workload measures.

This change in attitude and use of PCAM seems partly to be due to changes in NYPD management style, resources, and personnel. With the advent of the NYPD's COMPSTAT anticrime management system, which uses detailed, retrospective crime statistics to hold precinct commanders directly accountable for performance, dispatch delays are viewed as the result of individual precinct commanders' decisions on how to allocate their officers across competing assignments: responding to 911 calls, investigative work, addressing quality-of-life crimes, and special anticrime assignments; as well as how well they manage these officers. Moreover, there is now a relatively large patrol force available to deal with the stream of 911 calls, particularly those involving crimes in progress. The implicit assumption of the NYPD senior management seems to be that each precinct has enough resources to keep dispatch delays within desired standards and it is up to the precinct commander to determine how to do it (without the help of a PCAM model). This is in contrast to the late 1970s and early 1980s when, because of financial pressures, the size of the patrol force was repeatedly reduced while crime rates were very high. Finally, none of the current OMAP personnel has an OR/MS background, and their level of understanding of the model and its capabilities is generally limited. Of course, now it would be relatively simple to provide each precinct with its own laptop computer-based PCAM model, particularly because OMAP will soon be providing the precincts with the data needed to run it. Several precinct commanders, after learning of PCAM's capabilities during a Columbia-based

educational program for senior officers of the NYPD, have expressed a strong interest in having this tool.

Meanwhile, the director of the NYPD's OMAP told us that PCAM was still seen as a valuable tool and that he would like an update of the PCAM computer code to "improve efficiency" and to take advantage of the kind of advanced mapping capabilities that are at the heart of COMPSTAT to improve patrol allocation and perhaps dispatch decisions. He also expressed a desire for a study on the effect of changing the current, sometimes "irrational," precinct boundaries, which can separate two sides of the same street into two different precincts. Another study on his wish list would look at the potential improvements that may be derived by changing the current tour design, which still segments the day into three nonoverlapping time periods and uses a suboptimal "fourth platoon" as described above, even though a RAND study (Kolesar et al. 1975) showed this to be grossly suboptimal 25 years ago.

We found a somewhat similar situation at the FDNY Division of Management Analysis & Planning. A dog-eared copy of the RAND Fire Project book was on the shelf at the FDNY's Management Analysis & Planning Division and an updated version of the firehouse-siting model was still being used. In fact, the model played an important role in supporting Mayor Bloomberg's recent and controversial decision to close six firehouses in the face of an almost \$4 billion budget gap. This decision, like virtually all decisions to close firehouses, met with strong opposition from the affected communities and resulted in sustained demonstrations and hearings before the City Council (Colapinto 2003, McIntire 2003), where members questioned the model's validity. But in testimony before the City Council, the fire commissioner claimed that three months of postchange actual data confirmed the siting model's predictions that fire protection would not materially deteriorate. However, as with PCAM, our discussion revealed that none of the people running the siting model has an OR/MS background and, consequently, the way the model was used seemed suboptimal; the key response time measures were being calculated without weighting by incidence frequency. Somehow, over the years, the siting model's original capability to compute weighted average response times had been lost and no one understood the old computer code well enough to reinstitute this feature. Instead, the FDNY's analysts did ad hoc adjustments to the model's outputs to approximate the desired measures.

We also learned that the RAND-developed dynamic fire company relocation model was still in use, though it was now consulted less frequently than when first implemented during the years of high alarm rates in the late 1970s. But the algorithm played

an important role when the World Trade Center disaster on September 11, 2001, emptied most of Manhattan and much of Brooklyn of fire protection. More than 200 fire companies responded to the Twin Towers location, approximately half the entire city's complement, and the relocation algorithm was used to help rebalance the remaining resources in the city by relocating several dozen companies. The combination of the decision support provided by the algorithm, the sound judgment of the chiefs in charge, and the heroic efforts of the firefighters maintained fire protection at adequate levels throughout the rest of the city despite 9/11 being an otherwise average alarm rate day. A thorough picture of the FDNY's response to this disaster, coupled with numerous recommendations for improvements in communications and logistical planning for possible future disasters, is contained in the report by McKinsey & Company (McKinsey & Company 2002).

The specter of future terrorist events has prompted FDNY interest in augmenting the capabilities of the dynamic relocation algorithm to include real-time evaluations of alternative relocations suggested by chiefs or dispatchers in terms of expected response times in the affected area. The 9/11 experience has also prompted an application by the FDNY to secure funding from the Department of Homeland Security to develop a siting model for ambulances, which the department has been managing since they were given responsibility for Emergency Medical Systems several years ago.

Several observations emerged from our conversations with the two agencies. First, there is a continued appreciation and understanding of the need for computer models to support operational decision making. A dramatic example of this is COMPSTAT, the much recognized and imitated NYPD management innovation that has been credited as a major factor in the city's more than 60% reduction in crime over the last decade (but which currently incorporates no statistical or OR/MS models or analysis). Second, the development of information and communications technology makes possible approaches to deployment that were technologically or economically infeasible when emergency service modeling was at its height in the 1970s. Third, there are many challenges to the continued effective use of such a model after it is developed. These include changes in the management environment of the host agency, the need to maintain and update software and hardware, and personnel changes that result in the model being used by people who do not adequately understand how or why the model does what it does. Finally, and perhaps most importantly, the managers at both the NYPD and FDNY expressed the need and desire for new management science models to help with current issues of importance.

So What Does the Future Hold?

The advent of 9/11 and broad threats of domestic terrorism has given rise to an entirely different perspective on emergency responsiveness. Managers of police, fire, and ambulance systems, as well as mayors and governors, must now think about how to prepare and plan for catastrophic events that were previously unthinkable. As indicated by our conversation with the FDNY, 9/11 has created a new imperative for management science models and analyses to help design emergency systems and plans for responding to and minimizing the impact of terrorist attacks or other potentially large-scale emergencies. The reaction to the 9/11 attacks has also created a source of funding for research through the U.S. Department of Homeland Security. Some of this funding is allocated directly to municipalities and states through the Office of Domestic Preparedness. Other funding, through the Science and Technology division, is targeted for creating Centers of Excellence in universities to conduct multidisciplinary research to "enhance our ability to anticipate, prevent, respond to, and recover from terrorist attacks" (see Department of Homeland Security 2003). The first such center was created in November 2003 at the University of Southern California. Though this center is primarily focused on risk analysis related to the economic consequences of terrorist threats and events, it will contain a component addressing the development of models for emergency responsiveness. Future centers may include one that is focused exclusively on the latter. One might hope this could create an institution akin to the NYCRI with the potential to bring together a team of talented operations researchers focused on important problems of emergency responsiveness that has sufficient funding to sustain modeling and analysis over a period of years.

However, experience and history tell us that the potential for implementation and impact of new models and analysis will rest on several other factors as well. First, there is a need for well-defined client organizations that can act as full-fledged partners in the development of models to ensure their usefulness and actual implementation. Institutional and/or political leadership, as existed, for example, in the FDNY during the NYCRI era, is necessary to champion what may be controversial suggestions about new policies and practices. For example, the very fine work of Kaplan et al. (2002), which demonstrated the superiority of a mass vaccination program to minimize deaths resulting from a smallpox attack, met with significant political resistance from some public health officials and institutions, and implementation has been delayed. Also, as we learned from NYPD's experience with PCAM, the sustainability of any new operational support system depends heavily on the degree

to which there exists in the client agency an influential and technically competent planning organization with the ability to effectively use, maintain, and upgrade models over the years.

An additional obstacle to implementation of a new generation of emergency service decision support models dealing with high-impact emergencies is that resulting policies and plans will almost certainly involve coordination of various agencies and geographic regions. This greatly complicates both the politics and logistics of implementation. Another distinction is that unlike “routine” emergencies, which, by definition, occur regularly and for which there is ample data, there is no practical way to validate models of 9/11-type of events. This presents additional hurdles in establishing the credibility of models and in identifying unanticipated consequences and the nonobvious factors that need to be addressed.

Despite these challenges, the history of success in using modeling and analysis in emergency planning and responsiveness, as described in many *Management Science* articles as well as elsewhere, demonstrates that our field can and should play an important role in minimizing the impact of both routine and catastrophic emergencies in the future. We may now be facing an unprecedented opportunity to use our unique skills and experience to influence the collective welfare, and in doing so, recapture some of the energy and excitement that emanated from the origins of our field during World War II.

Acknowledgments

The authors wish to thank Andrea Allocca of the FDNY and Chief John Gerrish of the NYPD for meeting with us on current operations in their departments. Ed Blum, Jan Chaiken, Ken Chelst, Ed Kaplan, Dick Larson, Randy Hall, Art Swersey, Warren Walker, and Larry Wein offered valuable input and comments.

References

- Archibald, R. W., R. B. Hoffman. 1969. Introducing technological change in a bureaucratic structure. New York City–RAND Institute, P-4025, New York.
- Athanassopoulos, A. D. 1998. Decision support for target-based resource allocation of public services in multiunit and multi-level systems. *Management Sci.* **44**(2) 173–187.
- Blum, E. H. 1972. The New York City fire project. A. Drake, R. Keeney, P. Morse, eds. *Analysis of Public Systems*, Ch. 7. MIT Press, Cambridge, MA.
- Blumstein, A. 2002. Crime modeling. *Oper. Res.* **50**(1) 16–24.
- Brandeau, M., R. C. Larson. 1986. Extending and applying the hypercube queueing model to deploy ambulances in Boston. A. Swersey, E. Ignall, eds. *Delivery of Urban Services*, Invited Chapter. North-Holland, New York.
- Carter, G., E. Ignall. 1970. A simulation model of fire department operations. *IEEE Trans. Systems Sci. Cybernetics* **SCC-6**(4) 282–293.
- Carter, G., E. Ignall. 1975. Virtual measures: A variance reduction technique for simulation. *Management Sci.* **21**(6) 607–616.
- Carter, G., J. Rolph. 1973. New York City fire alarm prediction models: I. Box reported serious fires. Report R-1214-NYC, The RAND Corporation, Santa Monica, CA.
- Carter, G., J. Rolph. 1974. Empirical Bayes methods applied to estimating fire alarm probabilities. *J. Amer. Statist. Assoc.* **69**(348) 880–885.
- Carter, G., J. Chaiken, E. Ignall. 1972. Response areas for two emergency units. *Oper. Res.* **20**(3) 571–594.
- Chaiken, J. M. 1978. Transfer of emergency service deployment models to operating agencies. *Management Sci.* **24**(7) 719–731.
- Chaiken, J. M., P. Dormont. 1978a. A patrol car allocation model: Background. *Management Sci.* **24**(12) 1280–1290.
- Chaiken, J. M., P. Dormont. 1978b. A patrol car allocation model: Capabilities and algorithms. *Management Sci.* **24**(12) 1291–1300.
- Chaiken, J. M., E. Ignall. 1972. An extension of Erlang’s formulas which distinguishes individual servers. *J. Appl. Probab.* **9**(1) 192–197.
- Chaiken, J., W. Walker, P. Dormont. 1985. Patrol car allocation model: Executive summary. Report R-3087/3-NIJ, The RAND Corporation, Santa Monica, CA.
- Chaiken, J., E. Ignall, P. Kolesar, W. Walker. 1980. Response to the communication on RAND-HUD fire models. *Management Sci.* **26**(4) 422–432.
- Chelst, K. 1981. Deployment of one- vs. two-officer patrol units: A comparison of travel times. *Management Sci.* **27**(12) 213–230.
- Chelst, K. 1988. A public safety merger in Grosse Pointe Park, Michigan—A short and sweet study. *Interfaces* **18**(4) 1–11.
- Chelst, K. 1990. Queueing models for police-fire merger analysis. *Queueing Systems* **7** 101–124.
- Colapinto, J. 2003. The untouchables. *New York Times Magazine* (October 5) 54–59.
- Daley, R. 1973. *Target Blue*. Delacourt Press, New York.
- Department of Homeland Security. 2003. www.dhs.gov.
- Dickson, P. 1971. *Think Tanks*. Ballantine Books, New York, 249.
- Dormont, P., J. Hausner, W. Walker. 1975. Firehouse site evaluation model: Description and users manual. Report R-1618-2-HUD, The RAND Corporation, Santa Monica, CA.
- Drake, A., R. Keeney, P. Morse, eds. 1972. *Analysis of Public Systems*. MIT Press, Cambridge, MA.
- Fitzsimmons, J. A. 1973. A methodology for emergency ambulance deployment. *Management Sci.* **19**(6) 627–636.
- Green, L. 1980. A queueing system in which customers require a random number of servers. *Oper. Res.* **28** 1335–1346.
- Green, L. 1984. A multiple dispatch queueing model of police patrol operations. *Management Sci.* **30**(6) 653–664.
- Green, L., P. Kolesar. 1984a. A comparison of the multiple dispatch and $M/M/c$ priority queueing models of police patrol. *Management Sci.* **30**(6) 665–670.
- Green, L., P. Kolesar. 1984b. The feasibility of one-officer patrol in New York City. *Management Sci.* **30**(8) 964–981.
- Green, L., P. Kolesar. 1989. Testing the validity of a queueing model of police patrol. *Management Sci.* **35**(2) 127–148.
- Green, L., P. Kolesar, J. Soares. 2001. Improving the SIPP approach for staffing service systems that have cyclic demands. *Oper. Res.* **49**(4) 549–564.
- Greenberger, M., M. A. Crenson, B. L. Crissey. 1976. *Models in the Policy Process*. The Russell Sage Foundation, Basic Books, New York.
- Halpern, J. 1979. Fire loss reduction: Fire detectors vs. fire stations. *Management Sci.* **25**(11) 1082–1092.
- Hayes, F. 1972. From inside the system. A. Drake, R. Keeney, P. Morse, eds. *Analysis of Public Systems*, Ch. 1. MIT Press, Cambridge, MA.

- Henry, V. 2002. *The Compstat Paradigm: Management Accountability in Policing, Business and the Public Sector*. Looseleaf Law Publications, New York.
- Hitch, C. 1960. *The Economics of Defense in a Nuclear Age*. Harvard University Press, Cambridge, MA.
- Hogg, J. 1968. The siting of fire stations. *Oper. Res. Quart.* **19** 275–287.
- Hoos, I. 1972. *Systems Analysis in Public Policy: A Critique*. University of California Press, Berkeley, CA.
- Ignall, E., G. Carter, K. Rider. 1982. An algorithm for the initial dispatch of fire companies. *Management Sci.* **28**(4) 366–378.
- Ignall, E., P. Kolesar, A. Swersey, W. Walker, E. Blum, G. Carter, H. Bishop. 1975. Improving the deployment of New York City's fire companies. *Interfaces* **5**(2) 48–61.
- Kaplan, E. H., D. L. Craft, L. M. Wein. 2002. Emergency response to a smallpox attack. *Proc. National Acad. Sci.* **99** 10935–10940.
- Knapp, W. 1972. Summary and principal recommendations of the commission to investigate allegations of police corruption. The Fund for the City of New York, New York.
- Kolesar, P., E. H. Blum. 1973. Square root laws for fire engine response distances. *Management Sci.* **19**(12) 1368–1378.
- Kolesar, P., K. L. Rider. 1981. The fire department. C. Breshner, R. D. Horton, eds. *Setting Municipal Priorities 1982*, Ch. 9. Russell Sage Foundation, Basic Books, New York.
- Kolesar, P., A. Swersey. 1986. The deployment of urban emergency units: A survey. A. Swersey, E. Ignall, eds. *Delivery of Urban Services. Studies in the Management Sciences*, Vol. 22. North-Holland, New York, 87–120.
- Kolesar, P., W. E. Walker. 1974. An algorithm for the dynamic relocation of fire companies. *Oper. Res.* **22**(2) 249–274.
- Kolesar, P., W. E. Walker. 1975. A simulation model of police patrol operations. Report R-1625-NYC/HUD, The RAND Corporation, Santa Monica, CA.
- Kolesar, P., W. E. Walker, J. Hausner. 1975a. Determining the relation between fire engine travel times and travel distances in New York City companies. *Oper. Res.* **23**(4) 614–627.
- Kolesar, P., K. Rider, T. Crabill, W. Walker. 1975b. A queueing linear programming approach to scheduling police cars. *Oper. Res.* **23** 1045–1062.
- Koopman, B. O. 1956a. The theory of search part I: Kinematic bases. *Oper. Res.* **4** 324–346.
- Koopman, B. O. 1956b. The theory of search part II: Target detection. *Oper. Res.* **4** 503–531.
- Koopman, B. O. 1957. The theory of search part III: The optimum distribution of searching effect. *Oper. Res.* **5** 613–626.
- Larson, R. C. 1972. *Urban Police Patrol Analysis*. MIT Press, Cambridge, MA.
- Larson, R. C. 2001. Hypercube queueing model. S. I. Gass, C. M. Harris, eds. *Encyclopedia of Operations Research and Management Science*. Centennial Edition, Kluwer, Boston, MA, 373–377.
- Larson, R. C. 2002. Public sector operations research: A personal journey. *Oper. Res.* **50**(1) 135–145.
- Larson, R. C., T. Rich. 1987. Travel time analysis of New York City police patrol cars. *Interfaces* **17**(2) 15–20.
- Lawless, M. W. 1987. Institutionalization of management science innovation in police department. *Management Sci.* **33**(2) 244–252.
- MacDonald, H. 2001. Using Compstat against terror: NYPD crime busting system will work against Bin Laden. *New York Daily News* (November 4).
- Maltz, M. D. 1994. Operations research in studying crime and justice: Its history and accomplishments. S. M. Pollock, M. H. Rothkopf, A. Barnett, eds. *Handbooks in Operations Research and Management Science*, Vol. 6. North-Holland, Amsterdam, The Netherlands, 201–253.
- McIntire, M. 2003. City says closings slow fire response in six areas. *New York Times* (October 2) B4.
- McKinsey & Company. 2002. Increasing FDNY's preparedness. Report to the City of New York, New York, http://www.nyc.gov/html/fdny/html/mck_report/toc.html.
- Moore, M., G. Allison, T. Bates, J. Downing. 1975. The fourth platoon. Case C14-75-013, Kennedy School of Government, Harvard University, Cambridge, MA.
- Murphy, P. V. 1977. *Commissioner: A View from the Top of American Law Enforcement*. Simon & Shuster, New York.
- O'Hagan, J. 1973. Improving the deployment of fire fighting resources. *Fire J.* **67**(4) 42–46.
- Rajan, B., N. R. Mannur. 1990. Covering-location models for emergency situations that require multiple response units. *Management Sci.* **36**(1) 16–23.
- RAND Corporation. 1973. *RAND 25th Anniversary Volume*. The RAND Corporation, Santa Monica, CA.
- Ranzel, E. 1970. Garelick calls RAND study of city's police a failure. *New York Times* (October 7).
- Rider, K. L. 1976. A parametric model for the allocation of fire companies in New York City. *Management Sci.* **23**(2) 146–158.
- Sacks, S., S. Grief. 1994. Orlando magic. *OR/MS Today* **21**(1) 30–32.
- Savas, E. S. 1969. Simulation and cost-effectiveness analysis of New York's emergency ambulance service. *Management Sci.* **15**(12) B608–B627.
- Savas, E. S. 1973. The political properties of crystalline H₂O: Planning for snow emergencies in New York. *Management Sci.* **20** 137–145.
- Savas, E. S. 1978. On equity in providing public services. *Management Sci.* **24**(8) 800–808.
- Smith, D. 1969. *Report from Engine Company 82*. Warner Books, New York.
- Swersey, A. J. 1982. A Markovian decision model for deciding how many fire companies to dispatch. *Management Sci.* **28**(4) 352–365.
- Swersey, A. J. 1994. The deployment of police, fire and emergency medical units. S. M. Pollock, M. Rothkopf, A. Barnett, eds. *Handbooks in Operations Research and Management Science*, Vol. 6. North-Holland, New York, 151–190.
- Swoveland, C., D. Uyeno, I. Vertinsky, R. Vickson. 1973. Ambulance locations: A probabilistic enumeration approach. *Management Sci.* **29** 686–698.
- Szanton, P. 1972. Analysis and urban government. A. Drake, R. Kenney, P. Morse, eds. *Analysis of Public Systems*, Ch. 2. MIT Press, Cambridge, MA.
- Valinsky, D. 1955. A determination of the optimum location of fire-fighting units in New York City. *Oper. Res.* **4**(3) 494–512.
- Walker, W. 1975. Firehouse site evaluation model: Executive summary. Report R-1618-1-HUD, The RAND Corporation, Santa Monica, CA.
- Walker, W., J. Chaiken, E. Ignall, eds. 1979. *Fire Department Deployment Analysis: A Public Policy Analysis Case Study: The RAND Fire Project*. North-Holland, New York.
- Wallace, D., R. Wallace. 1999. *A Plague on Your House: How New York Burned Down and National Public Health Crumbled*. Verso Press, New York.
- Wallace, R., D. Wallace. 1980. RAND—HUD fire models. *Management Sci.* **26**(4) 418–422.
- Warner, G. 2003. One-officer patrol cars improve response time. *Buffalo News* (August 10).
- Wohlstetter, A. J., F. S. Hoffman, R. J. Lutz, H. S. Rowen. 1954. Selection and use of strategic air bases. Report R-266, The RAND Corporation, Santa Monica, CA.
- Worth, R. 2003. Times change, officials say, and so fire houses are closing. *New York Times* (April 11) D1.

Copyright 2004, by INFORMS, all rights reserved. Copyright of Management Science is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.