

BEING BUMPABLE¹

Richard I. Cook, MD
Cognitive Technologies Laboratory
Department of Anesthesia and Critical Care
University of Chicago

Analysis of decision making in complex domains often assumes a sharp distinction between narrow technical features and other organizational and institutional issues. In action, however, decision making by experts in medical domains engages organizational and technical issues together and simultaneously. I wish to argue (1) that much of practitioner expertise is focussed on and refined around the ways in which technical factors interact with organizational factors such as resource availability, procedural rules, and team issues and that (2) the rich, specialized argot that characterizes communication among practitioners within NDM domains includes terms that span technical *and* organizational factors. This paper presents an episode that occurred in a large teaching hospital. The episode reveals how practitioners cope with the demands of the socio-technical system of modern healthcare. It also leads to a broader suggestion of the ways in which NDM researchers might approach investigations of practitioner expertise in healthcare and, perhaps, other domains.

1. FIRST STORY:

A patient needed care in an intensive care unit after surgery. The patient was unable to be moved directly to the intensive care unit after surgery because there was no bed available in that unit. The patient was taken first to the recovery room and later transferred to the intensive care unit

2. THE SECOND STORY:

2.1 Background: "ICU"

An intensive care unit (ICU) is a place for patients whose precarious condition requires moment-to-moment attention by nurses. This sort of care requires close, often invasive (i.e. skin penetrating) monitoring equipment and specialized nursing skills. Within the hospital, patients typically come to the ICU from the operating room after a surgical procedure that leaves them unstable² or from the emergency room or, sometimes, from the *floor* after deterioration of their condition (e.g. after cardiac arrest). ICUs are sometimes classified as open or closed.

In an open ICU, the plan of care remains the province of the original physician. Here the ICU physicians (if there are any at all) are consultants. Patients in closed ICUs, in contrast, have the primary responsibility for their care transferred to the ICU physician. The responsibility for care includes the discharge from the ICU to the *floor* when the patient's condition is sufficiently improved (or sufficiently hopeless) that further ICU care is unwarranted. Whether open or closed there is considerable variation in the extent of participation by the original physician and the ICU physician in the care of a given patient. In the episode described below, the ICUs involved are open.

¹ To be published in the *Proceedings of the Fourth Conference on Naturalistic Decision Making, May 29-31, 1998*, The Airlie Conference Center, Warrenton, VA.

² The meaning of the word *unstable* is not precise but refers either to a condition where continuous interventions are required to maintain life or where there is a high perceived likelihood of abrupt change in condition that will warrant intervention. Examples of patients in ICU include patients requiring pharmacological support of blood pressure via intravenous infusions, those dependent on mechanical ventilation for gas exchange, or those with ongoing bleeding or fluid losses that require rapid replacement.

An ICU is an expensive place to receive care, usually an order of magnitude more costly than care on the *floor*. This reflects the costs of the one-to-one nurse-to-patient ratio of an ICU, and the extensive monitoring, radiographic and laboratory use that characterizes the ICU patient.

In addition to the formal ICU, there are also sometimes intermediate level of care (and cost) units, known as ‘stepdown’ units [because they are a ‘step down’ from the full fledged ICU] to which the less acutely ill may be transferred. Telemetry units may be considered stepdown units because the patients there have continuous radio transmission of their ECG to a central monitoring console. Typically the nurse:patient ratio in a stepdown unit is 1:2.

On the one hand, the variety of level-of-care units may be seen as providing physicians with a panoply of options for the care of individual patients. In practice, however, the distinctions create a particular type of physician burden. In such systems, the physician is responsible for evaluating the multiple dimensions of care required for individual patients and matching these with the individual unit’s characteristics. Because of competition for limited resources, the evaluation process is seldom a matter of simply selecting the perfect match. Usually it is a question of finding a match that satisfies the needs of the individual patient and also preserves other desirable characteristics of the workplace. The result is an overall increase in the complexity of the of the practitioner’s cognitive work that carries with it manifest opportunities for failure.

More broadly, the appropriate use of ICUs is a *cause celebre*. The debate centers on likely outcomes given disease processes and patient factors. Criteria for entry into and exit from ICUs are ill-defined, vary between institutions (and even within institutions!), and never crisp. Factors that weigh in the decision to admit to the ICU include details of the patient but also expectations of the quality of care that would be received in the ICU and the *floor* and the opportunity costs of committing the resources. Similarly the decision to discharge a patient from the ICU may include consideration of the expected future course of the patient, the costs of care, the consequences of inaccuracy in assessment³, and the availability of other resources such as stepdown or telemetry units.

Of particular interest in the surgical ICUs is the ability to anticipate some admissions. Certain surgical procedures⁴ usually require the patients to come to the intensive care unit for post-operative care; for these scheduled and, by definition, are *elective* cases, it is possible to predict the need for an ICU bed for post-operative care.

The number of patients that an ICU can handle is fixed for each work shift and fairly constant on the span of days. Under special circumstances some accommodation for an additional patient can be made but these are exceptional. Indeed, there are many adjustments made on a moment by moment basis to accommodate shifting workload.⁵ Many of these

³ Consider the effect of a *bounceback*. A *bounceback* is a patient who is discharged from the ICU but whose condition is insufficiently good to keep the patient on the ward and who returns to the ICU within hours or days. If, for example, the demand for ICU resources is acute and a patient is discharged from the ICU in order to use those ICU resources for another patient, a *bounceback* can create havoc. *Bounceback* also occurs at the floor level; an early discharge to home may result in a *bounceback* to the hospital. These events have specific costs and reimbursement consequences; *bouncebacks* after discharge to home may not be reimbursed as a new admission but rather as a continuation of the prior admission.

⁴ Examples are any procedure where post-operative mechanical ventilation will be required, very long procedures, and procedures conducted on patients already in an ICU.

⁵ For example, if there are several comparatively stable patients who require little attention and a sudden demand for care of a critically ill patient, the nurse in charge of the unit may assign a single nurse to several patients, freeing up

adaptations fall into the realm of unspoken conventions, the character of which is often specific to individual ICUs. Significant to the present episode, however, it is not generally possible to increase the number of patients that can be cared for in any ICU; in the short run, resources are fixed.

In the hospital in which the episode took place, there were eight ICUs of which five that were dedicated to the care of surgical patients. Of these five surgical ICUs, four were specialized (cardiac, burn, cancer, and neurosurgical services) and one was a general⁶ surgical unit. Although physically separate, for medical purposes, four of these ICUs (burn, cancer, neurosurgical, and general) were grouped together and act as a common pool of resources. [The cardiac surgical ICU was administratively separate and medically staffed by the cardiac surgeons and their residents who have no interaction with the other ICUs.]

The availability of beds in the ICUs is limited less by the physical number of beds than by the vicissitudes of nurse staffing; in this hospital there are more physical beds available than there are nurses to staff them. This is the case on the night of the episode and on most nights.

2.2 Background: floors

Most hospital patients reside not in ICUs but on floors⁷, each nominally associated with a particular surgical or medical service. There are neurosurgical floors, gynecological floors, etc. Patients coming out of a surgical ICU from a particular service are usually transferred to a corresponding *floor*. This is because the disease processes, patient characteristics, and postoperative problems that are most likely to be encountered are handled routinely on those floors. This is largely a matter of nursing expertise, which is refined over time to reflect

another nurse to attend to the new patient. But such adjustments have a ripple effect that extends through time. Accepting an additional patient for the current shift commits the unit to care for that patient during future shifts.

Temporary adaptations to handle brief workload surges are more common. For example, when a new patient is admitted to the ICU there is often a flurry of activity that requires ‘extra hands’. During this time, several nurses may cooperate, with two or three working on the individual patient while another one or two expand their scope of attention to encompass the patients for whom the other nurses were caring. To some extent the tempo of activity in the ICU is predictable: patients come from surgery at intervals, rounds occur at certain times, shift changes are regular.

Equally important but largely invisible to outsiders is the work *organization* done by nurses to cope with new admissions by shifting routine tasks associated with current ICU patients into relatively less demanding periods (system tailoring, Cook & Woods, 1996), batching processes (e.g. laboratory studies) for efficiency gains, and exploiting gaps in the work stream (e.g. the interval where an ICU cubicle is being cleaned) to perform other tasks (e.g. training, paper work, relief). To a limited extent, nurses can influence the work stream to create opportunities to handle overhead tasks effectively. Shift change, for example, is a period of intense activity the work may be planned so as to avoid overlapping discharge and shift change. Conversely, a discharge prior to shift change may be important to preparing to handle some anticipated future workload. The nature and predictable character of these demands tends to produce regular patterns of ICU activity such as the 03h30 drawing of blood for laboratory study that is necessary to have data available for 06h00 morning rounds by the physicians.

⁶ The adjective *general* when applied to surgery is in some cases inclusive and in others exclusive. A *general* surgeon is exclusively one with a certain type of training, usually five years post-medical school with training concentrating on surgery of the abdomen, breast, and thyroid. In the case of the *general* surgical ICU, at least in this episode, the term is more inclusive, referring to a physical location where patients from a variety of surgical specialties are housed. An exegesis on the term *general* as applied to surgery would trace the evolution of a complex interplay between growing domain knowledge, rapid advances in technology, and the internal political milieu of hospital staff organizations.

⁷ The term *floor* refers to a contiguous physical area that is also an organizational entity, e.g. “four north east” meaning the northeastern *nursing* unit of the fourth floor. The term “ward” was often used in the past, although it is no longer in vogue.

particular kinds of patients⁸. For example, patients who come from the neurosurgical unit require regular neurological checks to detect any change in neurological condition. Nurses on the neurosurgical floor perform many such *neuro checks*. For them these are routine while *neuro checks* are exceptional on other floors. Thus the strong preferences of surgeons for transferring their patients out of intensive care to “our *floor*” are quite rational.

2.3 Coping with complexity: parceling out beds by the bedmeister

In this hospital, the bed assignments for four units (burn, cancer, neurosurgical, and the general unit) are controlled by the *bedmeister*.⁹ The person playing this role is responsible for managing the ICU resources, matching the anticipated load of patients and planning for daily admissions and discharges. Such a role is essential to any functioning ICU but it is critical in the setting of multiple ICUs with overlapping capacities, where the individual ICU resource pattern is only one part of the overall hospital ICU resource situation. Because demand for ICU “beds”¹⁰ is only partly predictable, it is necessary to have a designated individual who is continuously aware of the resource situation and able to respond to new demands.

The *bedmeister* during the day is the intensive care unit fellow¹¹. At night two resident physicians are “on-call” to cover the four ICUs for general surgical patients. Because of the physical distribution of their respective ICUs, one is described as the “downstairs resident” (caring for patients in the general and cancer units) and one is the “upstairs resident” (caring for the burn unit and the general surgical portion of the neurosurgical unit). By convention, the downstairs resident is the nighttime *bedmeister*. Each of these residents is the primary physician responsible for the care of all patients in his or her units during the call period (typically from about 20h00 to 06h00). These residents are physically alone; specific issues may be referred to either the intensive care unit fellow or to the surgical service that is ultimately responsible for the care of the patient.

It is worth noting that the *bedmeister* role is a secondary one for the nighttime resident. The primary role is to care for the ten to thirty patients in the units, in itself quite often a demanding activity. Typically the on-call resident works through the day and then throughout the call night without any sleep but is the first to go home the following day, usually by 13h00 but occasionally as late as 18h00.

2.4 Cognitive tools: The Book

Elective surgical procedures that will require a post-operative ICU “bed” should not be started unless intensive care unit resources are assured. To allow for the orderly accommodation of these needs, a bound notebook (known as the Bed Book or just *The Book*) is kept in the general intensive care unit and used for “bed” reservations for these patients. A reservation can be made at any time but, because of the way elective cases are scheduled, most reservations are made 24 to 72 hours in advance. *The Book* is public and any surgical service can make a request

8 The physical grouping of patients also makes it easier for the physicians charged with their care, e.g. it makes the process of ‘rounds’ easier if all the patients associated with a particular service are co-located.

9 Bedmeister is a corruption of the German Bett (bed) plus Meister (master).

10 In this context, a “bed” does not mean the physical bed but rather the set of resources that are required for a patient; a “bed” is the smallest meaningful unit of ICU or ward resource. The term encompasses a collection of resources rather than a physical bed or patient cubicle.

11 A fellow is usually a physician who has finished residency training and is doing an additional year of specialty training, in this case in surgical intensive care.

at any time. The understood rule regarding “bed” requests is that they are honored on a first-come, first-served basis.

During the day, the ICU fellow and the nighttime *bedmeister* use *The Book* to anticipate the potential for future mismatch between the supply and demand for beds. For example, observing that tomorrow’s surgical schedule will overrun the available resources (given the current occupancy and ICU admissions expected today) will lead to a making the discharge of a specific number of patients today a high priority.

The incentives to discharge patients from the ICU to make room for new patients are not entirely aligned with the interests of the individual surgical service. Officially any *bed* that is emptied by discharging the current occupant from the ICU is returned to the *bed* pool. This means that the *bed* is available to the *bedmeister* for assignment to any surgical patient, not only patients from service that made it available. This reduces the incentive to a surgical service to discharge their patients from the ICU; the resources may simply be consumed by some other surgical service.¹² When a *bedcrunch* looms, however, the *bedmeister* can reduce the disincentive by agreeing to give the *bed* back to the discharging service for one of their new patients needing an ICU *bed*. For example, if the day begins with full ICUs and there are no discharges to the *floors* planned, the ICU fellow may refuse to allow elective surgical cases for a given service to proceed to surgery until that service has designated a *bumpable* patient (see below) from their service. Such arrangements tend to be *ad hoc* and politically sensitive for a variety of reasons.

The preceding has implied that ICU resources are homogeneous, i.e. that “a bed is a bed is a bed”. This is not completely true: ICU resources are not entirely interchangeable. The assignment of patients to ICU beds takes place under an evolving set of restrictions. Some of these are strongly medically based. For example, a septic (infected) patient should not be placed in the burn unit because of the increased susceptibility of the other patients in that unit to infection. Some other restrictions are more a matter of surgeon preference. Certain surgeons, for example, will not allow their patients to be sent to certain units because they believe the care in those units is not comparable to those in their favorite unit. Even if this opinion cannot be objectively supported, the realities of political power make such restrictions *de facto* rules of placement.

In most cases, the decision to transfer a patient to the *floor* is made by the primary surgical service during the morning rounds. A resident on this service writes the transfer orders and the patient is transferred to the appropriate floor as soon as there is a bed available on that floor. The ICU cubicle is cleaned and made ready for the next patient, a process that may take as much as an hour. This fits well with the normal daily activity: patients are discharged from the ICU to the *floor* in the morning, the patient is moved and the ICU bed area cleaned and prepared for a new admission while the patient who will occupy that bed is undergoing surgery. Later in the day, the new patient comes to the ICU cubicle and the process continues.

2.5 Preparing for demand > supply situations: identifying bumpable patients

Because the hospital is a trauma center and because patients sometimes become critically ill

12 The argument goes like this: Why should surgical service “A” discharge Mr. X who is, by the rather fuzzy criteria of judgement, is appropriately using ICU resources so that surgical service “B” can operate on Mr. Y? Operating on Mr. Y is elective and will only tie up the ICU resource for an even longer period of time than Mr. X is likely to remain. For service “A” to have a better chance of using the ICU for one of its patients it makes more sense to deny service “B” the “bed” currently occupied by Mr. X and to arrange to discharge Mr. X from the ICU only when another service “A” patient is ready to come to surgery that will lead to an ICU admission.

without warning, not all intensive care unit admissions can be planned. It sometimes happens that there is a demand for an ICU “bed” that cannot be immediately accommodated. In such situations it falls to the *bedmeister* to assess the patients in the ICUs and to select the least sick patient for transfer to a stepdown unit or the *floor*.

That this is possible at all is a function of the heterogeneous nature of ICU care and illness. Patients in the unit are in many different stages of illness and there are usually some who are nearly ready to be discharged from the ICU to the *floor*. These are candidates for earlier discharge to meet the immediate demands for ICU care for some other individual. In any event, discharge from an open ICU and transfer to the *floor* always requires the consent and active participation of the primary surgical service responsible for the patient.

When the ICUs are near capacity, the end of the day routine includes an informal conference between the intensive care unit fellow and the night's *bedmeister* to identify "bumpable" patients. A *bumpable* patient is one who is a candidate for transfer if an additional bed is required for a new patient. The advantage of identifying *bumpable* patients is that resolving new demands for ICU care is less likely to cause a major disruption and may be handled as routine. The notion of *bumpable* implicitly involves degraded levels of operations and the process of identifying *bumpable* patients is one means of planning for coping that characterizes many real world environments.

2.6 Son of coping: building an ICU from scratch when no ICU bed is available

If the units are full *and* ICU care for a new patient coming from the operating room is required *and* there are no *bumpable* patients a potential crisis exists that may require relatively extreme measures. Among these are the use of the post anesthesia recovery room (PACU) or the operating room (OR) itself as an intensive care unit.

These approaches are expensive, disruptive, and severely limited. Using the PACU or OR as an ICU uses nurses not generally involved in ICU care to provide ICU-like resources. The normal adaptive ICU activity cannot encompass these settings (see note 4). The per hour cost of PACU and OR are on the order of 4 to 10 times as much as the normal ICU.

Because these areas are physically isolated from the ICU, policy requires that a surgical resident stay with such a patient continuously until the patient can be transferred to a genuine ICU. This means that the surgical resident is unavailable for other work, for example in the emergency room or on the *floor*.

In addition, consuming the OR or PACU for ICU-like use makes them unavailable for other work. At night, when there are a limited number of OR teams available to do cases, tying them up as ICU personnel means that the ability to handle new OR cases is diminished.

2.7 Piling Pelion on Ossa: workload and new demands for attention

At night, the *bedmeister* function is only one part of the "downstairs" resident's job. The need for the *bedmeister* function increases as the other work of the resident increases. When the ICUs are not full and there is a new demand for a “bed”, the *bedmeister's* job is easy: it is only a matter of deciding to which of four units the new patient should be sent. But in this situation, the *bedmeister qua* care giver is relatively untaxed; because the census is low and “beds” are available, the resident confronts comparatively low workload. In contrast, if the ICUs are full, the *bedmeister* function becomes more difficult and time consuming while at the same time the number of patients requiring care is maximal. Thus the effort demanded to resolve requests for ICU resources is highly correlated with the other demands for resident time and attention.

2.8 *The episode sequence*

The episode occurred in a month when the ICUs were continuously working at capacity; the demands for “beds” were intense. For several days including the night of the episode, the units were full. Beds were available for new patients only as old patients were discharged. The following is a simplified sequence of events.

- 1) Surgical service “A” fails to reserve a bed in the open intensive care unit for a surgical case scheduled for the late afternoon and the case is placed on “hold” because no beds are available.
- 2) The ICU fellow OK’s the beginning of the surgical case, knowing that there are one or two *bumpable* patients in the unit. The fellow instructs the *bedmeister* to identify and bump a patient.
- 3) The *bedmeister* identifies the least sick patient in the unit. This patient is a trauma victim who had been admitted 24 hours earlier and has had a relatively stable course. Because the ICU is ‘open’ the *bedmeister* calls the surgical service caring for the trauma patient and requests permission to transfer that patient to the regular *floor*. The *bedmeister* also asks the ICU clerk to call the admitting office (which clears requests for beds on the *floors*) and obtain a *floor* bed assignment.
- 4) The surgery service declines to transfer the trauma patient but instead recommends another post-surgical patient for discharge from the ICU and transfer to the *floor*. This patient has had relatively minor surgery but is undergoing a ROMI¹³ procedure and so must be transferred out not to a regular *floor* but to a telemetry bed for continuous ECG monitoring. The *bedmeister* notifies the ICU nurses and ICU clerk of this change in plan.
- 5) The patient selected for transfer has a pulmonary artery catheter. These catheters are permitted only in intensive care units and the catheter must be removed and replaced with an intravenous line. The *bedmeister* changes the pulmonary artery catheter to a triple lumen catheter in a sterile procedure.
- 6) Two hours pass. The *bedmeister* attends to other duties.
- 7) The surgical procedure is proceeding smoothly and the circulating nurse in the operating room calls to see if the ICU bed is ready. The *bedmeister* notes a new laboratory value showing that the patient to be bumped now has a blood sugar of 450 mg/dl (high) and blood and acetone in the urine. The *bedmeister* concludes that the patient is in diabetic ketoacidosis (DKA) and cannot be transferred to the *floor*. The *bedmeister* cancels the transfer and begins treatment of the DKA.
- 8) The *bedmeister* identifies another *bumpable* patient for transfer. This patient had a severe nosebleed the night prior and required multiple blood transfusions. The patient has not had any bleeding for 18 hours. The *bedmeister* contacts the surgical service caring for this patient and obtains permission to transfer this patient to the *floor*. The patient does not require continuous ECG monitoring and so cannot be transferred to the telemetry stepdown unit but must be transferred to a regular *floor*. The *bedmeister* asks the clerk to contact the admitting office to

¹³ ROMI (pronounced “roam-ee”) is an acronym for Rule Out Myocardial Infarction. The process involves observation and laboratory tests for twenty-four hours after the presumed point of infarction. Positive evidence for myocardial infarctions includes various changes in the electrocardiogram and an increase in the serum concentrations of certain enzymes found in the heart and released after acute heart injury. Typically a ROMI includes a sequence of three blood samples drawn eight hours apart and daily ECGs. Because potentially life threatening irregularities in heartbeat are known to occur in the period after a myocardial infarction, the patient undergoing a ROMI has continuous ECG monitoring until the infarction is ruled out. Patients at risk for myocardial infarction are routinely ROMI’d after surgery because the stress of surgery may provoke an infarction and because the anesthetic may mask the symptoms of infarction (typically, chest pain).

clear a *floor* bed for this patient. The *bedmeister* discusses the situation with the ICU nurse in charge. They elect to "hold on" to the telemetry bed (and not to notify the admitting office that the telemetry bed is now not required) for possible use in the event of another request for an ICU bed, assuming that the first patient will again be *bumpable* at that time.

- 9) The ICU nurses notify the operating room nurse that the ICU bed will not be available for at least 1 to 2 hours. The transfer to the *floor* will delay the bed availability.
- 10) A surgical resident is asked to come to the ICU to write transfer orders for the patient being bumped. The resident comes to the ICU and writes the transfer orders.
- 11) The OR nurse foresees that the surgical procedure will be finished well before the ICU bed is available and notifies the operating room desk nurse who is instructed to alert the post anesthesia care unit (PACU) nurses. The PACU nurses normally leave the hospital after the last ordinary (non-ICU bound) operating room case is completed; patients going to the ICU are transferred there directly without first going to the PACU. In this case, the PACU will be used as a temporary intensive care unit and the PACU nurses will need to stay in order to fill the gap between the end of the surgical procedure and the ICU bed becoming available.
- 12) Because the patient will require a mechanical ventilator, something not normally used in the PACU, the PACU nurses page the respiratory therapist for a ventilator for the patient.
- 13) Elsewhere in the hospital, a bed is assigned to the *bumped* patient from the ICU and nurses prepare to receive that patient. ICU nurses prepare to transfer that patient. Cleaning personnel prepare to clean the ICU location as soon as it is vacated.
- 14) In another operating room, an unrelated *cardiac* surgical procedure is nearly finished. The circulating nurse in *that* room calls the *cardiac* surgical ICU and requests delivery of an ICU bed¹⁴ to the room. The cardiac surgical ICU nurses call for a "patient care technician" (a low skill hospital employee whose job includes transporting equipment) to move the bed from the cardiac surgical ICU to the operating room corridor.
- 15) A patient care technician picks up and delivers an ICU bed not to the area outside the cardiac surgical operating room but rather around the corner outside the operating room in which the episode case is proceeding. The technician is a relatively new employee and fails to recognize that the bed is placed in a non-standard position, closer to the room in which the episode case is going on than to the adjacent cardiac surgical suite where the patient for whom the bed is intended is undergoing heart surgery.
- 16) A nurse in the episode case operating room sees the ICU bed in the hallway outside the operating room. Because ICU beds are brought directly from the ICU cubicle to which they will return, this is interpreted by the nurse as a signal that the ICU is now prepared to receive the episode patient. The nurse concludes that the ICU bed is available and notifies the PACU nurses that they do not need to remain because the patient from the episode case can be taken directly to the ICU.
- 17) The PACU nurses notify the respiratory therapist on-call that the ventilator is no longer required.
- 18) The ventilator is removed from the recovery room. The PACU nurses prepare to leave the hospital.
- 19) The episode case is completed and the preparations are made to move the patient onto the ICU bed.

¹⁴ The reference here is to a physical bed, not to the resources associated with the argot term *bed*. Physical ICU beds are used for transfer of patients going to the ICU. This eliminates the need to transfer the patient from the operating room table to a gurney and then from a gurney to the ICU bed. ICU beds are distinctive because they have electrical controls that permit the bed to be raised and lowered, the head adjusted, and so on and are prepared with special padding in the ICU prior to being used for patients. They are much larger than transport gurneys, which are used for patients being moved about the hospital.

- 20) The *bedmeister* stops by the operating room to see how the case is going. The circulating nurse for the operating room comments that the patient is ready for transport to the ICU. The *bedmeister* tells the circulating nurse that the ICU bed is not yet available and that the plan to transport the patient to the recovery room has not been changed. The circulating nurse and the *bedmeister* exchange harsh words [!].
- 21) The *bedmeister* goes to the PACU and intercepts the PACU nurses as they are preparing to leave. They are informed of the fact that the original plan remains in force and ordered begin to prepare for the transport of the patient to them. More harsh words are exchanged [!].
- 22) The episode case patient's transfer is put on hold in the operating room pending the delivery of a ventilator to the recovery room, thus turning the operating room into a temporary ICU.
- 23) A ventilator is delivered to the PACU.
- 24) The episode case patient is transported to the PACU. Report of the case is given to the PACU nurses.
- 25) The *floor* nurse notifies the ICU nurses that a bed is now available for the *bumpable* patient. This patient is transferred to the floor.
- 26) The ICU cubicle is cleaned and prepared for the episode case patient.
- 27) The episode case patient is transported from the PACU to the ICU. Report of the case is given by the PACU nurses to the ICU nurses.
- 28) The next morning the *bedmeister* describes the sequence of events to the administrative head nurse of the ICUs. The nurse identifies the root cause of the difficulty as one of "human error" on the part of the patient care technician who delivered the ICU bed to the wrong hallway location.

2.9 The sequel: consequences of the failure episode:

The episode patient suffered no direct harm. A partial list of the consequences of this failure includes the following:

1. Extra operating room time and the time of two nurses and three physicians (surgical resident, anesthesiology resident, anesthesiology attending staff).
2. Extra recovery room time including the time of nurses and respiratory therapists; double setup of a ventilator.
3. Multiple exchanges of responsibility for the patient's care with resulting information loss about the conduct of the case and the condition of the patient.
4. Time and effort to schedule and reschedule beds. The loss of telemetry resources from the larger pool for one evening.
5. Extra time spent by the *bedmeister* away from the patients in the ICU.

3. Some observations on the episode

The administrative nurse's attribution of the failure to "human error" on the part of the patient care technician is remarkable. The focus on the most proximate flaw in a long series of flaws and faults as the "cause" is the usual circumstance in complex system failures (Woods, et al., 1994). Interestingly, the event occasioned no deeper reflection on the part of the administrator. No thought was given to the nature of ICU resource use or scheduling, the roles of various parties in resolving resource conflicts, the impact of operating a hospital at or near its theoretical capacity, the potential consequences of the extra activities required of the many people in the ICU and in the *floor* to cope with these circumstances, or the implications of the "hiding" of the (unused) telemetry bed. *What is remarkable about this episode is that the events*

were normal and even mundane.

The sequence appears almost Byzantine in complexity. Yet even this telling of the story is a pared down account derived from pages of notes. In fact, I argue that all accidents in medicine have similarly complicated stories that might be told but generally are not. The decision to tell a longer or shorter version is arbitrary. Reading and understanding this episode, even simplified as it is, is a demanding task. More significantly, the detailed recounting of episodes requires both technical knowledge (e.g. what does diabetic ketoacidosis mean?, what is a pulmonary artery catheter?) and organizational knowledge (e.g. how are ICU assignments made?) for comprehension. Comprehension also requires understanding of the local factors and roles (e.g. the bedmeister) that may differ substantially between hospitals and in a given hospital over time. As the story gets longer, the background knowledge required becomes ever more complicated. The present episode demonstrates the great difficulty in communicating succinctly the details of the sorts of events that permeate everyday practice. It also demonstrates the complex interconnectedness between domain knowledge and practice that is often unregarded or poorly qualified during investigations of accidents.

The episode also calls into question our commonplace views of medical knowledge as the ground substance of clinical decision making. There is little in the description about the underpinnings of medicine and nearly nothing about classical “medical decision making”. While all of these components are present in the work and the participants could describe them if questioned, they remain submerged within the activity of the parties involved. This is in keeping with the inner stories of many accidents (e.g. railway or shipping accidents) where the participants’ activities and expressions focus not so much on what (naïve) researchers regard as narrow domain details as on a broad set of pragmatic issues regarding the system in operation. The domain specific knowledge that outsiders consider crucial (e.g. what constitutes an unstable patient, what medical details are relevant in determining whether a patient is *bumpable*), are embedded in the workaday world knowledge of the participants. For practitioners, the narrow domain details are distinctly secondary to the main issues with which they must deal.

Practitioner argot is often regarded as simple shorthand for specific technical domain language. But the suggestion here is that argot performs a more complex function than simply standing in, one for one, for technical terms. The practitioner language provides useful aggregates of domain knowledge. These aggregates refer to factors within *and across* the levels of goal-means abstraction hierarchy (Rasmussen, 1986). Being *bumpable* means having temporarily taken on a certain characteristic that has individual and larger group meanings. Medical stability, the threat of future clinical events, the ease with which these may be detected, the consequences of failing to detect them, all play into the assignment of this state. But being *bumpable* is more; it depends not only on inherent characteristics of the patient but on the states of *other* patients, the needs of *other* units, the *other* available resources. Being *bumpable* involves comparisons, contrasts, and expectations. It is not a static circumscribed technical term, like “coronary artery disease” or “diabetic ketoacidosis”, but a fluid, dynamic, complex, time limited, assessment – characteristics that it shares with the domain. It subsumes a variety of local factors, local issues, local limits and local capabilities. *Bumpable* is not ephemeral or insubstantial; it has real meaning that is subject to negotiation between practitioners and throughout the organization. The argot is a powerful tool used to express the relationships between higher and lower abstraction levels, neither just technical nor just organizational but rather both. That practitioners develop and employ such argot indicates that their work is not divided between technical and organizational elements but about the interplay between these.

The widespread use of terms like *bumpable* indicates that the properties to which they refer are enduring characteristics of the domain.

Bumpable has at least one other important quality in this context, namely its forward projection. *Bumpable* refers to a present state with future consequences. It is inherently about planning for a potential future, of anticipating and preparing to cope with unknown and unknowable future events. It references a hypothetical future time in which the practitioners will have to deal with a degraded system. It shows one means by which practitioners reach *forward* into a hypothetical, conflicted, demanding future. By declaring patients *bumpable*, practitioners envision a future world which will demand that they act in ways that push the system towards saturation and, potentially, towards overt failure.

Galison (1997) has pointed out the presence of specialized “trading languages” in the high energy physics community. These allow individuals from different disciplines (e.g. refrigerating engineers, experimental physicists, theorists) to exchange information about the interacting characteristics of their work. The argot of intensive care practitioners is an instance of this. *Bumpable* allows nurses, surgeons, administrators, and cleaning people to exchange information about a particular patient, to declare intentions, to argue about the directed meanings of technical and organizational factors, to consider the future pace and character of work, and to foresee hazards. *Noto bene* that argot terms do not and cannot have a universal meaning for everyone. For a nurse *bumpable* implies, perhaps, preparation of papers and communications for transfer. For the nurses’ aide it implies a future empty bed that will need to be made ready for a patient. For the resident it implies a resource to be consumed at a future time. The term provides both discipline specific meanings and simultaneously a shared representation that allows the work to be organized and communication to proceed.

The very presence of these terms is a clue to the difficult, demanding, and hazardous locales within the domain of practice. The notion of *bumpable* is directed at the margin of practice rather than at the median of the distribution. Because this system typically runs at or near capacity, it is the marginal events and circumstances that are the focus of attention, that threaten to saturate resources, that generate disputes, rancor and confusion, and that elicit expertise. There are other terms (e.g. *bounceback*, *bedcrunch*) with similar characteristics. Argot thus plays a central role in the planning to react and the reacting by planning that characterizes NDM domains and this plan/reaction is directed not at the system’s center of mass but the edges where failure and success hang in the balance.

The suggestion here is that practitioner argot points directly to the critical concerns of NDM researchers. Because they span the distance from technical to organizational, these terms capture the essence of what is going on and therefore are at the locus of the decisions/actions of practitioners. To explore what it means to be *bumpable*, what a *bedmeister* actually does, what a *bedcrunch* entails, what *The Book* actually does, is explore the core of success, failure, and expertise.

References

- Cook R, Woods D. Adapting to New Technology in the Operating Room. **Human Factors** 38(4), 593-613.
- Galison, P (1997). **Image and Logic**. Chicago: University of Chicago Press.
- Rasmussen, J (1986). **Information Processing and Human-Computer Interaction**. NY: Elsevier.

Wood D, Johannesen L, Cook R, Sarter N (1994). **Behind Human Error**. Wright-Patterson AFB, Dayton, OH: CSERIAC.