NECSI

NECSI Technical Report 2004-09-01

September 2004

System Care: Multiscale analysis of Medical Errors—Eliminating errors and improving organizational capabilities

Yaneer Bar-Yam, Ph.D.

New England Complex Systems Institute

The systems basis for medical errors is widely acknowledged. How to improve organizations to perform more effectively is less well understood. Complex systems concepts can be used to analyze the origins of medical errors as well as methods for changing the system to reduce their frequency or eliminate them. The key to this analysis is recognizing that the source of a particular error does not lie in a sequence of events that occurred in a particular circumstance, but rather in the set of possibilities that can occur and the ability of the organization to match the set of necessary possibilities with a corresponding set of appropriate actions. An analysis of the space of possibilities enables an identification of the "weak links" in organizational structures. Application to the problem of prescription fulfillment suggests that a likely source of the widespread errors in this context is the convergence of communication from diverse specialists to the pharmacists or administering nurse. Specific recommendations follow from this for organizational improvement either by increasing the redundancy of this communication channel, or by developing care teams that include medical professionals associated with local groups of patients. Effective and ineffective use of automation is discussed.

Article Series Introduction:

Rapid improvement in the quality of care has become a high priority for the healthcare system. Despite the expansion of medical knowledge, the use of increasingly sophisticated technology, and the high level of physician training, measures of the quality of care, return on investment [1] and the incidence of medical errors [2] depict a severely under performing system.

This is the second of four articles that address the healthcare system using recent fundamental advances in complex systems research.[3] The central analytic tool that will be used in these papers is Multiscale Analysis [4,5], which identifies the functional effectiveness of a system by considering the set of actions a system can perform at different scales.

The four articles in this series identify the interplay between individual and system capability for particular organizational forms and demonstrate both why and how the current difficulties in the healthcare system exist. They also provide direct recommendations as to how effectiveness can be dramatically improved, while keeping within the financial constraints of the system. These articles explain why system structure and behavior rather than individual competence or negligence should be the center of attention for quality improvement and error reduction strategies. Beyond this recognition, they provide specific strategies for such improvement efforts. The four articles address in turn:

(1) *The role of the financial and organizational structure of the healthcare system in inducing resistance to efficiency improvement and ineffectiveness—pointing to the need to separate types of tasks, simple and repetitive tasks from complex and unique ones, and have different organizational forms address each type of task to enable both efficiency and effectiveness.*
(2) *The role of complexity in the inability of existing organizational structures to reliably perform increasingly complex tasks in individual care—pointing to the need for a local team-based structure, which distributes but integrates observation, decision making and actions so as to enable tasks that are much more complex than any one individual can perform.*
(3) *The ineffectiveness of prescribed protocols, planning and assignment of responsible individuals to oversee complex tasks—pointing to the need for organizational learning, superceding individual training, as a means for creating effective teams that can perform complex tasks.*
(4) *The limitations of technology in addressing high complexity tasks—pointing to the need for appropriate technology and the recognition of the proper place for its use in the efficient and effective healthcare and public health system.*

In this, the second article in the series, we discuss a set of approaches to reducing medical errors with specific application to the problem of prescription fulfillment errors. The discussion focuses on five natural approaches to error reduction: feedback correction, eliminating steps, redundancy, automation, and reducing the local complexity of the task. All of the approaches may be useful under some circumstances, however, the analysis suggests that for the current prescription error problems, increasing the redundancy should provide the best strategy for achieving immediate improvement. This conclusion arises from the existence of converging communication channels between specialist physicians and those who fulfill their instructions. A broader view on the organizational structure suggests that there is an increasing need for team based medical practice. Teams provide a significant increase in ability to deal with complexity over the current strategy of increasing specialization.

The following sections are, in brief: 1) the growing recognition of the significance of medical errors, 2) the problem of prescription errors, 3) the space of possibilities analysis methodology, 4) approaches to reducing medical errors, 5) reducing the local complexity of tasks, 6) evaluation of the improvements that can be expected, 7) conclusions.

Medical Errors

In recent years, the health care industry has grappled with an increasing awareness of its own fallibility. A 1999 Institute of Medicine (IOM) report[2] announced that preventable medical errors of all kinds are killing between 44,000 and 98,000 people per year—more than the number of deaths due to automobile accidents or breast cancer. While the methods of counting and accounting for medical errors are disputed, the problem has become acute no matter whose numbers you trust. The dangers associated with receiving medical care have become a growing concern for the American public. Dramatic examples of medical errors—often fatal—appear regularly on the front pages of newspapers and the covers of magazines. The public's perception of medical errors is often dominated by a scapegoat mentality that prompts reporters and readers to assign unambiguous blame to a particular individual, procedure, or device.

The need for widespread improvement has been recognized, but it's not always clear what kind of framework can help health care providers understand how these errors come about. The IOM has emphasized that the key to reducing medical errors is an understanding that they are "systems related" and not attributable to individual negligence. Recognizing that the errors come from system design is a positive step, but it doesn't actually tell you how to improve the system to prevent the errors from occurring. Before we can understand how to improve the system, though, we need to look at what is needed to provide effective and error-free medical care. In this paper, we will construct a portrait of a medical system that offers highly effective and extremely specific individual care. The third paper will discuss in greater detail key strategies for improvement of the system to achieve such behavior.

Prescriptions and the problem of providing medication to patients.

There are many aspects of patient care where medical errors arise. We will discuss the example of the prescription error problem, one of the most common and extensively studied forms of error. The lessons we will learn from this example can also be applied to other areas where problems arise.

Providing medication is an important—and complex—service that the medical care industry provides. One way of understanding the complexity of a task is to count the number of possible options. How complex is the task of drug prescription and delivery? Today there are about 15,000 registered drug names in the United States. Supplying the right medication for a patient, then, means making sure that he/she receives the right one of those 15,000 possibilities. But that's not all; not only are there numerous drugs available, but there are many possible dosages—both quantity and timing—and methods of administering them. With all of these different parameters, imagine all of the possibilities, many of them potentially harmful. Nurses recognize this complexity and use a five "rights" mantra —the right patient, the right drug, the right time, the right dose, and the right route. Given a high-complexity task, where there are many wrong outcomes for each right outcome, errors are likely to occur. Conversely, if many errors are taking place, it's very likely that there's a high complexity task that isn't being dealt with effectively by the existing system.[3] The problem with the system for providing medication is that for many years it hadn't been revised to accommodate the increased complexity of its task. Today there are many efforts to improve the system, however, to make new systems work well, it is important to understand why the old system is failing.

For example, let's imagine how the traditional system might work for inpatient medications. The doctor scribbles the prescription on the patient's medical chart, possibly using certain well-established abbreviations. Then, a hospital employee copies it from the chart. The copy is taken to the pharmacy, where a pharmacist reads and fills the prescription. He gives the medication to a hospital employee (perhaps the same one, perhaps not), who then transports it to the appropriate area of the hospital, where a nurse finally administers the medication to the patient.

Let's start by examining one segment of the process: the doctor scribbling the prescription on a piece of paper. Theoretically, a doctor might have 15,000 possible medications to choose from when writing a prescription. One high-profile aspect of this proliferation of choices is name confusion. Take the example of these two drugs:

Celebrex and Cerebyx. Celebrex is a prescription medication that provides pain relief from arthritis. Cerebyx, on the other hand, is an anticonvulsant prescribed for the treatment of seizures. Name confusion has led to mistreatment of patients and this pair is only one example of the many pairs of similarly named drugs that have caused confusion in the writing or filling of prescriptions. Some further examples are Lamictal (an anticonvulsant used to treat bipolar disorder) and Lamisil (an antifungal drug); Zyrtec (an antihistamine) and Zantac (an ulcer drug); Sarafem (an antidepressant) and Serophene (a fertility drug).

Prescription errors also occur in specifying the correct dosage. For example, in a highly publicized case in Washington, DC,[6] a surgeon wrote a prescription for ".5 milligrams" (not "0.5") of morphine for a nine-month-old baby, to be administered by a nurse after a series of operations. The unit clerk transcribed this number as "5 milligrams," without a zero or a decimal point and the medication was dispensed in that amount. The nurse tending the child followed the order, and due to the erroneous dosage—ten times the intended amount—the child died.

Given this account of what happened, we might blame any one of the people involved in the prescription-filling process. We could argue that the doctor made the crucial error in leaving off the extra "0" before the decimal, making the number more open to possible misinterpretation. Or we might insist that the clerk's misreading and misfiling of the prescription was responsible. We could also contend that the nurse who administered the medication should have recognized that the dosage was too high for a small child. In the flurry of attention that followed this case, all of these hypotheses for who was "at fault" were proposed.

Space of possibilities

This account of the events leading up to the tragic error, however, leaves out the most important information of all: the space of possibilities for each step. Each of the individuals involved in this case had a distinct set of possible choices in the actions that he or she took. The set of possibilities for each task determined the likelihood for error. Without understanding the space of possibilities, we simply cannot evaluate the system to determine where the errors are coming from.

For example, what if morphine were *only* administered in the amount of 0.5 milligrams, to any kind of patient? If this were the case, the pharmacist and the nurse should have known that there's never, *ever* an instance in which 5 milligrams of the drug should be dispensed. On the other hand, if morphine were usually administered at 5 milligrams, then more responsibility might lie with the doctor, who should have been more careful to emphasize that this was an exceptional case by adding the extra zero and perhaps making the decimal point more visible.

Figure 4.1 presents a graphical illustration of the space of possibilities for this problem. The dots represent the set of possible outcomes for the decision-making process in administering medication to a patient. Each dot (in a multidimensional space) is a possibility that under some circumstances could be correct; each possibility is defined by the type of medication, dosage, route, patient, and time of administration. For a given situation, we want one and only one of these possibilities to occur.

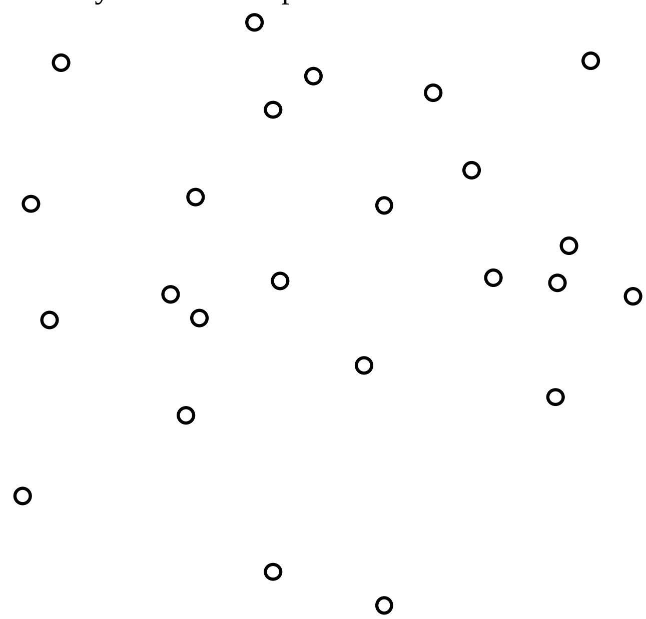

Figure 4.1: The space of possibilities: Each dot represents a possible valid decision.

Ideally, when a doctor writes a prescription, he will record information that corresponds to a complete description of one of these possibilities. Then, through the process of filling the prescription, the correct choice should be made. Now, loosely speaking, the complexity of a system is the amount of information needed to determine which of these dots is the one that has happened (or should happen). One measure of this is the length of that description—the number of letters, perhaps, used to record it. Therefore, the complexity of a particular prescription can be measured by the length of the description the doctor has written down.

What happens when errors occur? If a doctor miswrites a letter, or a pharmacist misreads a letter, the prescription no longer describes exactly the correct possibility. In Figure 4.2 the rings around the dots represent these errors of perception. Accidentally switching one letter for another in writing the name of the medication, for example, would correspond to moving out to the first ring around the dot. Another error would place us in the second ring. Because some of the error distances around the dots overlap in this space of possibilities, a small number of errors can take us from one distinct possibility to another.

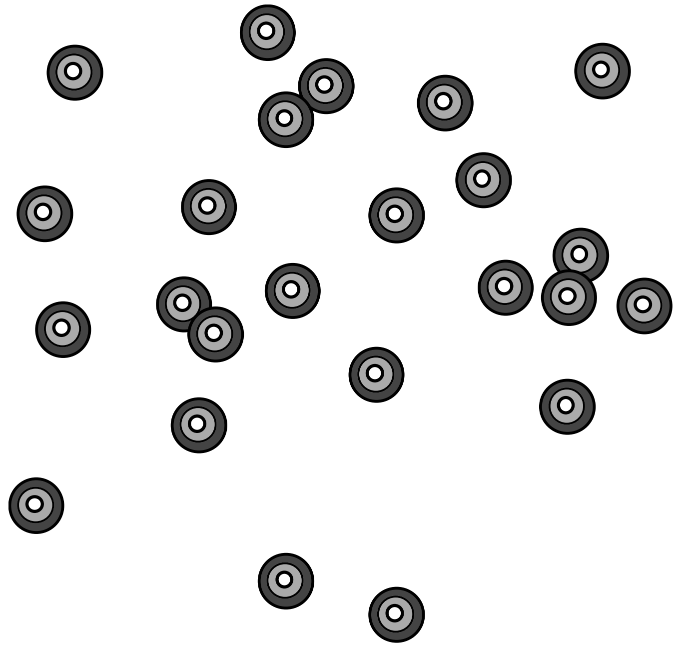

Figure 4.2: The rings represent the effects of errors. As long as the rings don't overlap the intended (right) decision can be inferred.

Suppose a doctor were writing a prescription for Cerebyx, but he makes three errors in his handwriting: he accidentally writes an "l" instead of "r," and an "e" instead of "y," and then inserts an "r" after the "b." With those three errors he's now written "Celebrex," a completely different drug. His errors have moved him from specifying one possibility to specifying another. If he only makes one of those errors (and writes "Celebyx," for example), the prescription will lie somewhere between possibilities. At this point, it's ambiguous which possibility is called for—did the doctor mean to write "Celebrex" or "Cerebyx"?

From this discussion, we can see that it's crucial to understand the structure of the space of possibilities. If there were no drug with a name very close to Cerebyx, then one or two wrong letters might not make such a difference. (For example, if there were only two drugs available on the entire market, Cerebyx and Prozac, then accidentally writing "Celebyx" would still unambiguously point to Cerebyx.) If all medications were administered to patients at a dosage of 0.5 units, then even the dropping of the zero from the prescription would not lead us ambiguously close to any other possibility—because there would be no other possibilities for dosage. The further away the dots are—or the fewer dots there are at all—the less likely you are to make enough errors to cross the space between them.

Error prevention and correction

In response to the medical error problem, many organizations have produced recommendations, proposing a variety of procedural, organizational, and technological changes that hospitals, clinics, and pharmacies can carry out to reduce errors. Many of these recommendations are quite reasonable. However, unless you have a sound understanding of the system you're trying to change, it's difficult to understand which changes will really help. It is also often hard to motivate people to make a change without being able to explain why it should work and even how well it should work. The key is understanding the fundamental role of complexity and scale and this is what will be used to analyze the proposed changes.

There are five quite natural approaches to reducing medical errors: feedback correction, eliminating steps, redundancy, automation, and reducing the local complexity of the task. The first four when appropriately used and effectively implemented can ensure that a decision that has been made is actually carried out. The last (reduction of local complexity) also has another use: reducing the likelihood of decision-making errors. Each of these approaches will be discussed below.

In talking about these methods of error reduction, it's important to be clear about one thing: we're not talking about subtle errors of judgment in the actual medical decision regarding what treatment is necessary. The errors we're talking about are obvious differences between what should be done (as decided by the physician) and what is actually done.

The diagram in Figure 4.3 illustrates the process we are concerned about. One decision maker, D (usually a doctor seeing a patient), makes the decision about the right action to perform. This decision is communicated through a series of intermediaries who carry out the intended treatment. The methods of error correction address the problem of deviations from the desired track—resulting in a different treatment than was intended by the doctor.

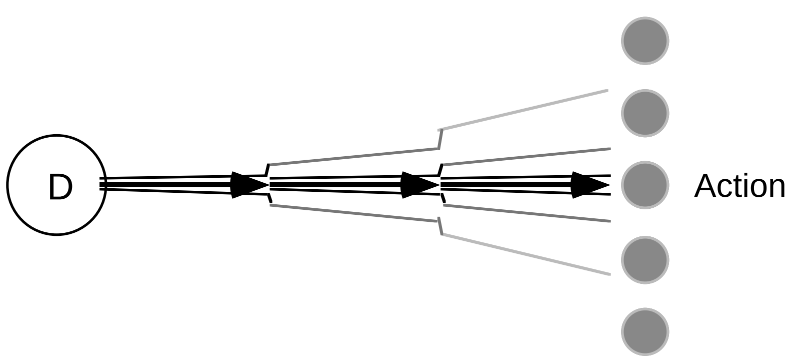

Figure 4.3: A decision followed by several communication steps, with possible errors, and then action.

The model of medical practice described in this diagram doesn't always apply, but it is helpful to think about this as a first step toward considering some of the key issues associated with medical errors. Still, we should be aware of the assumptions it requires. First of all, we assume that the only decision-maker in this process is the doctor. Under this assumption all other health professionals, be they nurses, technicians, or pharmacists, simply carry out the practical details of a decision which has already been made. The pharmacist simply translates the information from the doctor (the prescription) into the medication, and the nurse just administers that prescribed medication. The second and more subtle assumption is that the pharmacist receives only instructions from one kind of doctor (not one doctor but rather one kind of doctor).

Neither of these assumptions is necessarily true. Later in the paper we'll discuss their limitations, and what happens when they don't match what's really going on. Starting with a simplified picture will enable us to introduce the basic strategies for preventing errors. When we add more real-world complications we can see how they affect the usefulness of each of these strategies.

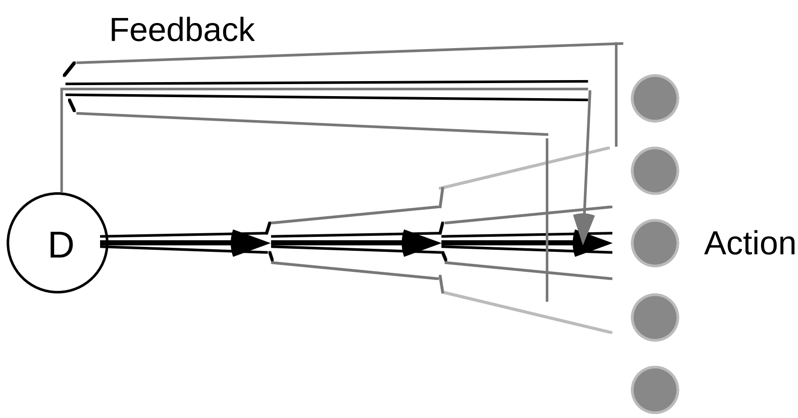

Figure 4.4: Adding a second path of information allows for feedback correction, but setting up and coordinating the extra path often creates its own opportunities for failure and surely is a lot more effort.

Feedback correction — Check once, check again

If errors are occurring in your system, one way to remove them is to put checking procedures in place to catch errors that have already occurred. For drug prescription and delivery, this kind of "feedback correction" involves double-checking the prescription at the end, or possibly at various stages in the process.

The most direct way would be for the doctor herself to check the medication before it's administered to the patient. Ideally, here's how this would work. The doctor writes a prescription in the hospital, which then follows the ordinary routes—it's taken by the hospital employee to the pharmacy, where the prescription is filled. Then, the prescription returns to the doctor, who checks to make sure that the medication is what she meant to prescribe in the first place. If so, then she gives the green light and everything is as it should be. If it's not, then the error has been caught and the prescription is sent through the same process again.

This scenario might seem sort of unrealistic. Doctors have a lot to do and it is unreasonable to expect that they will double-check every medical procedure in the hospital. Furthermore, coordinating the doctor's schedule to be at the right place and time would be ridiculously hard. However, there are more feasible ways to carry out this approach. For example, if the particular medication needed is accessible nearby, this kind of double-checking by the doctor might be possible. This happens in limited ways already for some medications that are "on-hand" on the hospital floor, emergency room or even doctor's office.

The existing hospital procedure also has a more general double-checking procedure. The prescription written by the doctor stays on the patient's chart, which is kept with the patient. A copy of the prescription is made, which is the copy that is taken to the pharmacist. Once the prescription is brought back to the patient, it can be checked against the original that stayed with the patient. Of course, as we've found out from the case in Washington DC, the additional copying adds a step that might itself introduce errors. Making sure that the initial copy is a good one requires care and automatic methods like using carbon paper, a photocopier or a fax machine, may ironically add other opportunities for problems (poor copy quality, malfunctioning equipment, and a need for adequate supplies, repairs and backup systems).

With all double checking procedures, we create two paths for the information instead of just one (see Figure 4.4). One copy of the prescription is sent to the pharmacist, where it is filled. The other route doesn't directly involve the medicine; it's only function is to keep an accurate record of the information in the prescription used to determine the medication. Once the medicine is obtained, it is double-checked against the other copy to make sure that the medication is the one that was originally prescribed. This type of double-checking procedure would catch errors that occur between the act of the doctor's writing the prescription and the actual administration of the medication.

However, there are several problems with this approach. First, this approach creates additional steps were errors can be introduced. Creating an extra path for information requires at least two additional acts, one at the beginning when the information splits into two paths, and one at the end when they are checked against each other. Moreover, this approach would not catch errors in the first step of the process, when the physician actually writes the prescription. If the prescription is written with an error and we duplicate the prescription, the error now exists in both copies. Because the first step of in the process plays a special role, we will pay particular attention to it in the discussion of each of the strategies for reducing errors. To solve this problem, we must consider the act of writing out the prescription as the first step in the communication channel and find some way to duplicate that step. Before we discuss how this can be done, however, let's consider the possibility of removing unnecessary steps.

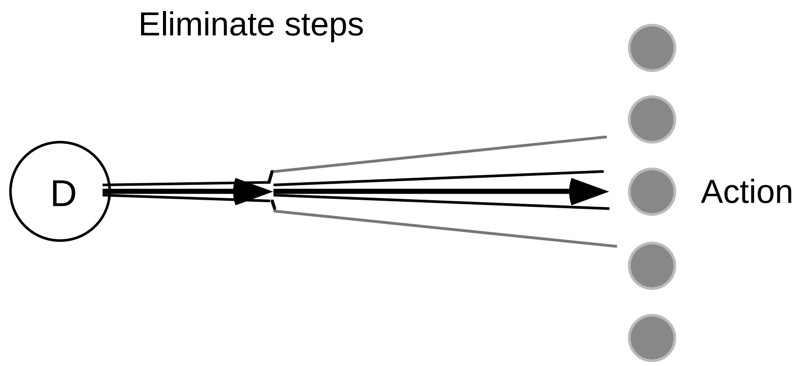

Figure 4.5: Removing unnecessary steps (here there is one less step than in Figure 4.3) reducing the possibility of error when those steps are likely to cause errors.

Removing unnecessary steps

Another important approach to reducing errors is to eliminate steps that might introduce errors (Figure 4.5). In the Washington DC case, for example, if the prescription had not been copied, perhaps the decimal point would have been noticed and the patient would have been given the right prescription.

If a current procedure contains unnecessary steps, removing the unnecessary steps reduces the likelihood of errors in the original process, which is better than having to eliminate errors once they have been made. This is also a good approach for reducing the amount of time needed to complete a process. For example, recently in a different hospital in Washington DC, the number of steps required to receive the results of a blood test was reduced from 8 steps involving 7 people and taking about 60 minutes to 3 steps involving 3 people and taking only 3 minutes.[7] This change in procedure was achieved by placing a small blood-testing facility right in the middle of the emergency room. With this arrangement, the person who draws the blood can immediately take the sample to the testing location, rather than having to send the blood sample to a different part of the hospital, saving a great deal time and making the overall process much more efficient.

However, there are also some problems that may occur when we try to eliminate unnecessary steps. First, the extra steps may be needed for other purposes. For example, if we want to have feedback checks as discussed in the last section, then extra steps are necessary. While reducing the number of steps reduces the likelihood of error, if we eliminate key checks we may actually end up increasing the number of errors. Evaluating the tradeoff (between adding steps that allow checks and removing them so they don't add more error) requires careful thought. Eliminating communication steps is also not possible when you want people to work together sequentially so that the task can be distributed among them. This often is the case when specialists or special equipment are necessary for part of the process.

Moreover, the approach of eliminating steps, like feedback checking, does not affect the very first step: the writing of the prescription. We can't eliminate that step (unless the doctor administers the medication) and preventing errors in the first step is important to ensuring that errors will not occur.

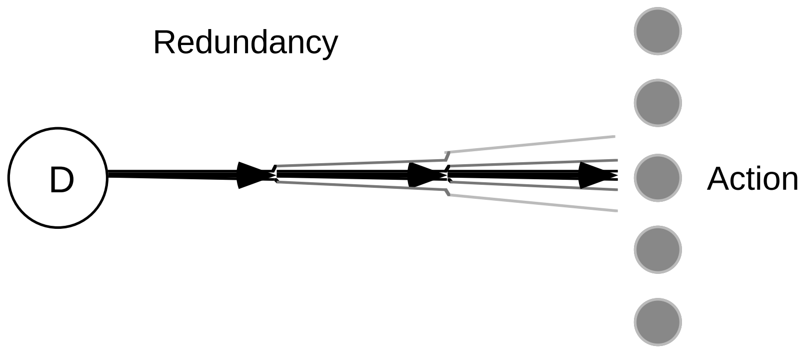

Figure 4.6: Redundancy adds more information, thus reducing the likelihood of errors at each step (including the first one), causing the spread of possibilities to decrease as shown.

Redundancy

The third approach to preventing errors uses redundancy. To create redundancy, one starts with more information at the outset of a procedure. The key to reducing the chance of error is to obtain more information from the physician at the start of the process. This information then follows along the entire route of the process so that everybody on the way can check to make sure that what they are doing is correct, thereby reducing the overall likelihood of errors (Figure 4.6). Having a lot of extra information could be burdensome, but it turns out that even just a little more is enough to reduce errors dramatically.

To implement this approach for medications the doctor would include twice as much information about the desired medication on the prescription. Using more words than the minimum necessary to specify which possibility is intended, provides a redundancy that can help to eliminate errors. If we consider a prescription that would be implemented correctly without the extra information, adding additional words gives nothing new and seems unnecessary and a waste of time. However, when there are errors and the extra information leads to a correct choice, it makes all the difference in the world.

For example, all doctors could write down on prescriptions both the generic and trade names of a drug, or they might write the name of the drug and the condition that it's being prescribed for (the indication). Any kind of additional information that could be used to identify the drug needed could be required on the prescription—including, even, the shape or color of the packaging This information is redundant, but it serves as a check on the other, standard form of description.

If you always write the drug name and the condition, then accidentally writing “Celebex; Seizures” instead of “Cerebyx; Seizures” would still indicate very clearly that you’re prescribing the anticonvulsant Cerebyx, and not the pain medication Celebrex. By increasing the amount of information you’re giving about the prescribed medication, you make the space of possibilities more and more dispersed. In effect this increases the distance between the dots in the space of possibilities shown in Figure 4.1, because the number of errors that would be necessary to go from one medication to another is quite large when there is more information. As the dots move farther and farther away from one another (with added redundancy of description), errors are less likely to matter. If the dots are farther away, then even four errors won’t lead to any dangerous ambiguity.

This is also the reason why physicians are advised to write 0.5 and not just .5 when writing prescriptions. The former has enough information to be interpreted correctly most of the time while the latter is more prone to error because the redundancy is low.

Like feedback correction, adding redundancy to a procedure means adding time to the doctor’s task. However, because of this redundancy, for example, having a prescription with both the name of the drug and the condition on it, the pharmacist (or the nurse, or the patient) may notice and be able to correct errors before administering the medication. In terms of complexity and scale, this is the same process as having the doctor re-approve the medication before it’s administered. In both of these procedures, you’re doubling the information that comes from the doctor so that the two sets of information can be checked against one another. With redundancy you double the information at once and the two sets of information, physically attached to each other, can be checked against each other at every step. With feedback the two sets of information are kept separate, and you check them against each other at a specific time later on. These two approaches are not exactly the same in the way they avoid errors, but they are close.

One of the crucial advantages to the redundancy approach is that it reduces the impact of errors in the very first step, the writing down of the prescription. No matter how the physician communicates the information the first time, including in automated ways that we will discuss next, the issue of making sure that this step is done well is crucial and redundancy can help.

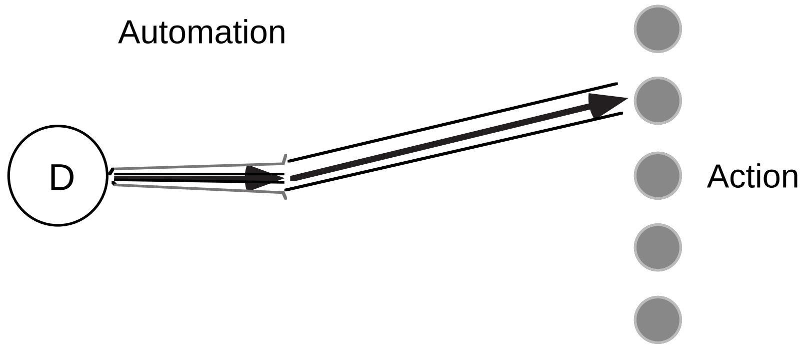

Figure 4.7: Automation reduces the likelihood of random mistakes occurring in the steps that have been automated, but may introduce problems at the starting point and through hard to find errors in implementation (bugs).

Automation

Automation involves identifying processes and chains of events that don’t require complex decisions and making them more efficient by introducing computers and communication technology. This often also reduces the number of people or steps involved to eliminate handoffs or communications that may cause error.

Why are computers helpful in reducing errors? To start with, it’s because they are less complex than people. Introducing a human being into a process produces the potential for change, because human beings are so complex. People are better at making subtle complex decisions than they are at automatic (rote) execution of simple tasks. For a given situation, there are potentially thousands of possibilities for what a person might choose to do. A computer, on the other hand, is not nearly as complex as a human being. It carries out repetitive, simple logic very reliably.

Automation is one of the most reflexively suggested methods of error reduction, but it’s not always the answer. It is interesting that computers are proposed as the best way to avoid human errors when the most commonly used computers frequently crash. There are two key problems with automation: correct implementation and an effective user interface. If the system is not implemented correctly, the system will make many errors. This illustrated in figure 4.7 by showing the process moving in the wrong direction. Since people believe that automation is the answer to solving problems, they will usually blame the programmer for implementation errors rather than the approach of using automation itself. If the user interface is not done correctly, there will be many errors that occur at the first step of the process, when the equipment is instructed what to do. When such an error happens people generally blame the person who entered the information rather than the user interface and do not think of blaming the strategy of using automation itself. As the complexity of tasks increases, it becomes increasingly difficult to make sure these two are done correctly. Indeed, sometimes it is much more difficult than having people do them correctly.

More generally, to really make proper use of automation, it’s crucial to understand what it is useful for. Automation will make things more *efficient* when the execution of a task can be uniquely and fully specified without any further decisions being made. It can also reduce error due to the elimination of unnecessary intermediate steps. For example, if, as we've been assuming, the doctor’s prescription uniquely specifies what should be administered to the patient, then the process of fulfilling the doctor’s decision involves no further decisions after the prescription is produced and automation can help. A hospital might set up a direct channel of information from the doctor to the pharmacy. The doctor would fill out an electronic prescription entry form that would immediately be sent directly to the pharmacy and printed out. The pharmacist would also do no copying whatsoever and would simply fill the prescription as printed out. We might go even further and install automated dispensing units, at least for common drugs, that receive the prescription and dispense the medication without any human intervention at all. Automated dispensing systems are already being implemented in some places.

Let’s look more carefully at the first step, the doctor writing the prescription. One part of this step in which automation may seem to be helpful but is only if it is done correctly, is in providing an additional immediate feedback check at the time of writing the prescription. A physician generally checks the prescription immediately after writing it. She reads it to check that it is clearly written as far as she is concerned. An electronic entry system can be designed to enhance this immediate check by having a typeface version of the written prescription or by automatically showing additional information like the standard medical indication for that prescription. While this may be helpful, some words of caution are worthwhile. This automated process seems like the same as the case where the physician wrote additional information. It isn’t. The additional information is not coming from the physician, it is only being verified. Verification has a lower level of reliability because it requires much less information from the physician. A physician is less likely to misspell a prescription and write the wrong indication (that happens also to correspond to the same drug that the spelling mistake gives), than to “blindly” approve an incorrect indication suggested by the electronic entry system. Thus, even if automation is used, it is better to have the entry system require the physician to enter both the medication and the condition. The key is to realize that the process of information transfer from the physician to the communication channel should not be made efficient. Despite the great desire to make it easier, the key to avoiding errors is to require more information from the physician as opposed to less. Once electronic entry is completed, feedback checking at the time of administration will be easier. Feedback involves sending information into two channels that contain identical data, which can then be checked against each other. One set of information is transferred to the pharmacist and translated into the medication, which then physically passes to the patient. The other channel is the feedback channel, which will simply contain an electronic version of the doctor's original prescription. Since the electronic version can be sent around automatically in any number of copies, the feedback process is simpler and, if the equipment is reliable, more reliable.

Feedback checking could be further augmented by having the computer read the package and do the comparison of medication with prescription (rather than the person administering). A few pharmacies and hospitals have adopted barcoded drug selection procedures, in which a paper prescription includes a computer-generated barcode that can be deciphered automatically at the pharmacy before dispensing, and even at the patient's bedside directly before administration. According to the FDA's recent regulations (February 2004) most prescription drugs, and over-the-counter drugs frequently used in hospitals, will be required to bear a bar code uniquely identifying the drug, its strength, and its dosage form. Checking this information with a barcode reader at the patient bedside—especially in conjunction with barcoded patient bracelets—could catch errors involving the wrong medication, dosage, timing, or patient.

There are many other useful ways for introducing automation into a system—electronic medical records, hand-held wireless computers for bedside use—but recognizing what they can improve and what they might not be able to do is important. Electronic medical records are important in enabling easy retrieval or sharing of information. However, among other issues, ensuring that the most important information is brought to the attention of the person who needs it, is not easy to guarantee. Hand-held computers can help in various tasks including checking medications. Some of the potential for improving the system and the required care in execution is described above. In each case the choice of what to automate and the quality of implementation are crucial to the development of effective systems. Since the existing systems have been developed and refined over many years, it will be difficult for new systems to introduce improvement unless great care is taken.

Two decision makers

Up until now we have been assuming that there is only one decision maker in the system—the doctor. However, it is definitely not that simple. For example, pharmacists also are responsible for making decisions—they're not just following instructions. Pharmacists are often responsible for determining whether multiple drugs prescribed to the same patient are incompatible, that is identifying harmful drug interactions before they occur, or making substitutions of one drug for another.

How does the complexity of the pharmacist's decision-making affect the possible solutions we've discussed so far? Some of the suggested improvements may not work as well or even at all in this case, while others survive unscathed.

Feedback correction now has a problem. Since the pharmacist can make drug substitutions, there may be good reasons that the prescription is not the same as the drugs that are administered. A simple feedback checking process will not work. The checking process, whether manual or automated, has to be able to figure out whether a substitution is OK or the result of an error. Either the person who is doing the checking or an automated system that performs checking must recognize which substitutions are reasonable and which are not.

Adding redundancy in the prescription still works. It improves the communication channel to the pharmacist without interfering with the pharmacist's decision making and allows him to modify the prescription if appropriate. In this case, writing both the medication and the indication seems like a really good solution.

Eliminating intermediate steps between the physician and the pharmacist, or between the pharmacist and the act of administration may still be helpful, but it cannot eliminate the pharmacist involvement. The same is true of automating steps in the process. The automation should not interfere with the decisions that are made by the pharmacist. Indeed, understanding the decision making role of the pharmacist is a key issue in determining whether or what steps to automate.

The many-to-one communication channel problem

Let's expand our view of the system one more step to observe that there are many different physicians sending prescriptions to the same pharmacist. Because of specialization there are many different types of physicians, and each specialty will tend to have its own set of most commonly prescribed drugs. While there are some drugs that a neurologist (a specialist in nervous system disorders) and a rheumatologist (a specialist in arthritis) might both prescribe, there are many others that are particular to each specialty.

From the neurologist's perspective, there's little to worry about when writing a routine prescription for Cerebyx for an epileptic patient. If there were another drug with a similar name that the neurologist tended to prescribe, he might naturally be more careful to identify clearly which he meant. But neurologists don't often prescribe Celebrex and from his point of view the communication path to the pharmacist may seem good enough. Similarly, a rheumatologist whose prescriptions happen to pass through the same pharmacist may see little likelihood for confusion in her own writing of a Celebrex prescription.

But this is not at all how the pharmacist sees it! The problem is that while the physicians have no need to think of both possibilities, the pharmacist is faced with both regularly and confusion is very likely indeed. More generally, there are many more possibilities on the pharmacist's side of the communication than on the physician's side. Having many different people communicating with one person places a very high demand on the communication channel at the far end. This is why it is not really enough for the physician to consider his own handwriting and ask himself if it is clear enough. He must consider what the pharmacist sees—how many possibilities the pharmacist has to distinguish among—to appreciate what he really needs to be clear about.

If one of the major problems lies with an unbalanced communication channel why hasn’t this view received more attention? The answer is quite simple: differences in perceived authority between physicians and pharmacists. Physicians are assumed to be more important and to have more power than pharmacists. Because of this, even if a pharmacist is uncertain about what drug the physician is prescribing, he might be reluctant to call up the physician to double check. Power is a key aspect of how roles are designed in an organization, and the weak points in organizational effectiveness are often determined by how power is perceived.

Differences in power generally are a way of shifting burdens from the powerful to the less powerful. Among these burdens is that of complexity. When physicians are powerful they can shift some of the complexity of their tasks onto others. If pharmacists become more powerful they could shift some of their complexity back to the physicians. This might not be a good idea if physicians have to respond to many calls from pharmacists at a time when they are already overburdened. Understanding which is better can only come from a more careful understanding of how the complexity of tasks is distributed through the system.

Now that we understand the problem and why it exists, how can we change the system to address the problem? Insisting on the physician’s writing the prescription in two ways, for example, the drug and the condition (indication), seems like the most direct solution. Or every prescription could be labeled with the specialty of the physician writing it. This would be a weaker, but possibly sufficient way to include extra information. Marking the specialty of the physician that gives a particular prescription might be done in a partially automatic way, for example by including it on prescription pads or through using an electronic identification system. This way the system can distinguish between what different physicians may write, without imposing special rules to make it work.[8]

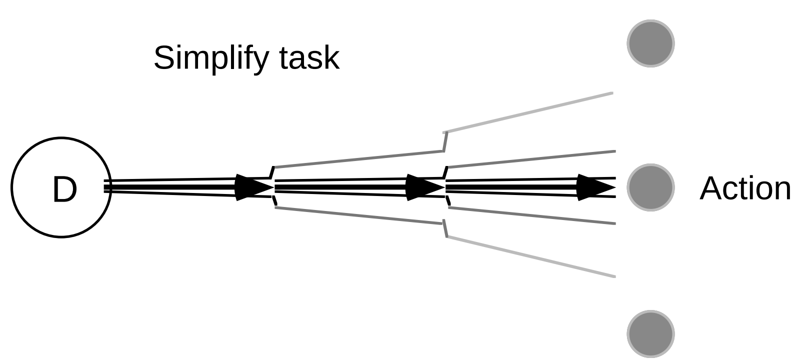

Figure 4.8: Simplifying the task reduces errors by making fewer options, shown here by reducing the number of dots representing valid actions. If fewer options are needed to perform the task, then this can work. Otherwise, specialization has to be part of this approach.

The next step in our discussion will change its focus. While we will still consider the communication channel, we can also use the next approach to address wider issues in the organization. This is important because the communication channel itself is not the only reason that errors occur. The wider view is necessary to address many other sources of error and thus to solve systems-related medical error problems.

Reducing local complexity

The use of feedback, redundancy and automation can help with the problems associated with communication channels. They reduce the demand on, increase the capacity of, or reduce the error rate in the communication channel. However, in many cases the source of errors may be the complexity of the tasks the individuals within medical system have to perform. The number of possibilities that medical practitioners face at every decision may be too large. A crucial method of error correction is to reduce the number of options available at any given step in the process (Figure 4.8). By limiting the set of possibilities that could be chosen, you reduce local complexity and thus reduce the likelihood of mistakes.

There are two ways to reduce the complexity that a person has to deal with. The first is to simply reduce the number of actions that the entire system can execute, and thus the number of possibilities the individual has to deal with. The second is to divide up the many possibilities among multiple individuals. Whatever changes are made to reduce local complexity, it's important to assess whether the overall task still has sufficient complexity to be effective. This is the crux of the problem of organizational effectiveness: you want your system to perform high-complexity tasks, but with individual local tasks that are simple enough that errors are unlikely to occur.

Reducing unnecessary possibilities: Standardization

The elimination of possibilities starts from the recognition that in practice, you don't always need *all* possibilities that might in principle be used. We see this process of stripping away unnecessary possibilities in many forms of standardization. For example, in the past pharmacists were responsible for mixing ingredients to produce medications in various forms (liquid solutions, ointments, powders, tablets and capsules). Today, however, pharmacists do much less mixing and packaging of the drugs, which usually come prepackaged in standard forms. Also, nowadays the dosage for many drugs is the same for all adults, often administered at fixed times twice a day. Drug packaging provides only a few options for how it can be administered. All of these changes reduce the set of possibilities tremendously. As long as the possibilities at your disposal correspond to the possibilities needed for treatment, this reduction in complexity is a very good idea.

In general, standards of practice lead to reduction of complexity. To the extent that we can be certain that the possibilities we are eliminating are absolutely unnecessary, this is great. However, when people develop standards they often consider only the "typical" or "average" case and create standards that do not apply to the space of all possible cases. Even in the case of standard adult drug doses there is the potential for problems: the same dose can have a very different effect depending on whether it’s administered to a football player or a jockey. This is the danger inherent in standardization: reducing complexity when it is needed for effective action.

Automation provides additional methods for standardization and constraints on the possibilities. For example, an automated system of electronic prescriptions could be used in quite reasonable ways to constrain the possibilities. The drugs could be organized according to condition being treated. Both the condition and the drug would have to be entered, in effect enforcing the redundancy recommendation given above. However, once the condition was entered, the set of medications that might be specified could be restricted. This means that the system automatically constrains the medication choices depending on the condition being treated. Alternatively, drug choices could be constrained by the specialty of the practice, or the name of the physician, or even the history of the physician’s pattern of prescriptions. With this kind of standardization, the doctor would select from a restricted number of choices. The automated system would use the information already entered to winnow the possibilities to choose from, reducing the possibility of error.

This kind of automated standardization would mean, for example, that a doctor prescribing pain medication for an arthritic patient would be unable to prescribe Cerebyx by accident instead of Celebrex if he has already specified that he’s treating arthritis. Such a system, well implemented, could be a reasonable automation of the process we described earlier of redundantly identifying the drug with the condition. However, it is important not to constrain the independence of the physician too much. The system must have procedures by which the doctor can override the standardized set of options; otherwise, the doctor’s limited choices might not allow exceptions necessary for specific cases.

One example where standardization does not appear to work is the conventional drug formulary system used by many health care organizations. Drug formularies are designed to limit the type of drugs that can be used. This was supposed to save money by limiting the prescriptions to lower cost drugs, when there were roughly equivalent lower cost and higher cost options available on the market. However, studies suggest that such plans have had the opposite effect,[9] increasing spending, while at the same time decreasing overall quality of care. Among the reasons for this outcome are the need for doctors to go through special administrative procedures to receive approvals of exceptions, and the use of “second-best” treatments that later required further medical care.

The pharmacist’s task and specialization

It is important to develop an understanding of task complexity to understand why solving the communication channel problem discussed earlier might not be sufficient for diminishing overall error levels. Let's take our best example of a method to fix the communication channel problem: writing both the indication and the drug. This approach seems like a very good solution for the communication channel problem and it might actually solve the prescription drug problems.

However there is a limit to this solution’s overall effectiveness. Consider the pharmacist who receives the prescription. People quite generally separate different types of information to different parts of their brain so that they can make composites that are the basis of creativity.[3, Chap. 2] While these composite states are the enabler of creativity, they are also a key source of errors. It is possible that a pharmacist would, therefore look at “Celebrex; Seizures” on a prescription but fail to notice the error. His brain may not see the incompatibility because of dissociation.

Right now the risk of this happening is not likely to be very high, but it’s important to recognize that this could become a major problem if the complexity of drug prescriptions reaches a high enough level. The complexity might increase as the number of names of drugs increases to the point that there will be enough combinations of drugs and conditions to create confusion. It is also important to realize that the dissociation we spoke about varies from individual to individual quite a bit. So it is possible to select the people who are naturally (or by effective habit of action) good at making sure that both the drug name and the indication are consistent with the drug given. If a person makes an error, then we could reasonably consider whether improved training is needed or that someone else would be better at the job.

Still, what can be expected even from very proficient people is restricted by the complexity limit of the individual. Once the necessary tasks surpass this limit, we need a different solution, one which assigns the tasks to multiple people rather than to a single individual. This is what happens in specialization, which can take many forms.

The first approach is to divert cases into separate channels. Using this approach you can limit how many kinds of cases a particular individual deals with, reducing the complexity of his task. Specialization is a very important and effective technique for complexity reduction. We'll understand its importance more clearly if we take a look at the usual medical routing system.

Figure 4.9 is a diagram of a standard medical routing system. This arrangement isn't universal (it does not apply to the emergency room), but it's still fairly typical. The white circle represents the general/family practitioner and the numerous black lines are the many patients who come to see him. These patients have a very wide range of conditions, and rather than treat them all himself he's referred them to specialists, the colored circles.

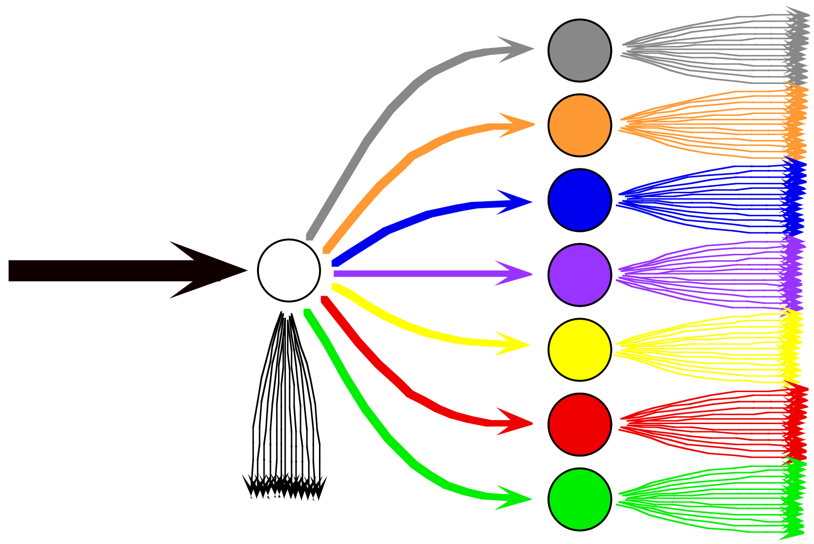

Figure 4.9: Medical routing system: The general/family practitioner receives all patients, makes some decisions and routes the others to specialists, who will be responsible for making further decisions.

The family practitioner thus deals initially with an extremely large variety of possible conditions. Her task, however, is limited to addressing directly a more limited set of conditions and routing (assigning) the rest of the cases to the specialists. The specialists don't have to deal with the same level of complexity as the family practitioner. Each specialist receives patients with a much smaller assortment of similar or related conditions. The specialists allow the primary doctor to forgo treating certain patients, so that the actual treatment of the patient happens at a much less complex level. This makes a lot of sense—you're separating the cases so that the set of cases that any one person has to address is less complex. Still the overall process has a much higher complexity, which is clearly necessary to address the individual case needs. This is the point of specialization.

However, this diagram does not completely represent the entire routing system. Where do these cases go after the specialist? They go (via the prescription) to the pharmacist, the rightmost circle in Figure 4.10. This circle might also represent the nurse administering the medications to the patients themselves or providing other aspects of the care.

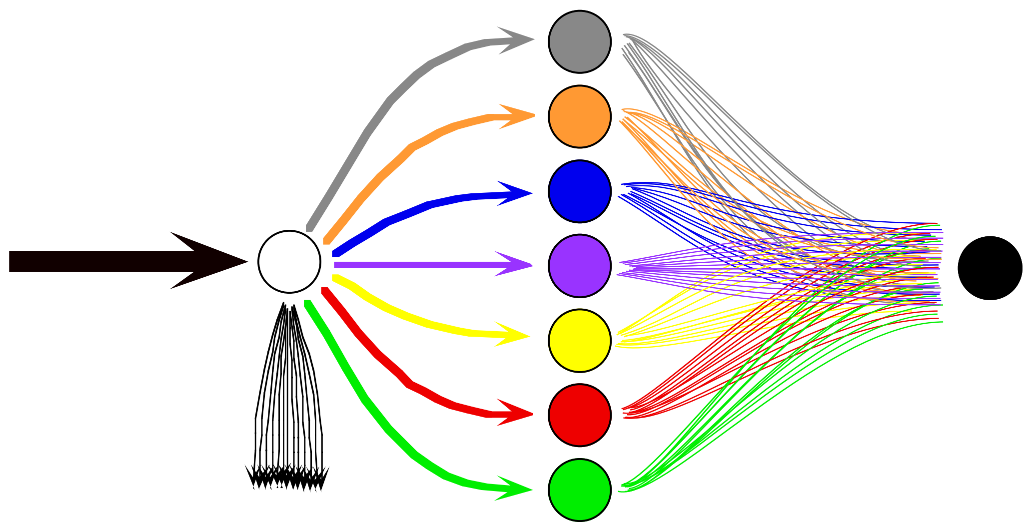

Figure 4.10: All specialists send patients to a pharmacist for medications (the general/family practitioner also does so; this is not shown in the figure).

What's wrong here? The cases have been separated because together they're too complex for one doctor to handle. Now, however, they're reunited again (either in the pharmacy or in the care of the administering nurse.) It's certainly not the case that all these patients now have similar needs—they still have vastly different treatment programs ahead of them. Of course, the full complexity of all of the cases does not fall on the pharmacist. There are many aspects of treatment aside from medication, and every condition does not require a distinct medication (consider the many different kinds of infection that are treated by the same antibiotic). Still, the routing system does reveal where problems are likely to arise. The architects of this system have applied complexity reduction to one part of the process, by referring patients to specialists, but have failed to do so at the other end of the channel, making it pretty obvious that the system's weakest point will be at the pharmacist/nurse, through whom all treatments must pass.

What's happened here is a good example of how systems adapt to increasing complexity. In general, as we learn more about how to treat medical conditions, the complexity of medical care increases because we learn how to effectively address a more highly varied set of cases using more specialized treatments. The community of physicians has addressed this increase in complexity by increasing the level of specialization, but the other parts of the system (like the pharmacy) have not found a natural way to adapt, so it's to be expected that problems will primarily arise at those points.

In order to make this system function more smoothly it is necessary to apply specialization to more than just one part of the process. From this diagram, the most obvious functions that should be specialized are the pharmacy and nursing care. There is already a limited degree of pharmacist specialization and nurse specialization. Nurses are specialized for emergency rooms, intensive care, anesthesia, and other areas. However, nursing specialization has been reduced in recent years with the cutting of costs, even as nurses' tasks have become more and more complex. The question is, how much specialization is appropriate? While it is clear that physicians have the most need for specialization, some degree of specialization of other professions in the tasks they perform is also likely to be necessary.

The importance of specialization can be found at many levels of organization. Specialization at higher levels of organization such as the care team or hospital would reduce the need for further specialization among the professionals who are working within that system. For example, we can consider the possibility of institution specialization, as found in children's hospitals, oncology (cancer) hospitals, trauma and burn centers. A pharmacy at an oncology hospital will be highly specialized for the very complex problems of drugs for cancer patients, for example. The existence of these specialized hospitals implies the importance of specialized knowledge reflecting the high complexity of care for patients in the categories (children) or with the conditions (cancer, trauma and burn) they treat. It also reflects the existence of a sufficient number of such patients to require a free standing institution. Developing specialized hospitals for every medical condition is not justifiable because the same circumstances do not apply, and because many patients experience multiple conditions.

At a lower level than hospital specialization, it is worthwhile to consider the strategy of forming teams. A team, consisting of doctor, nurse, and pharmacist (or a limited number of all three) can deal with the entire process of deciding what medication to use, filling the prescription, and administering the medication to the patient. If the unit is specialized to deal with certain types of cases, the number of distinct cases and possibilities that each individual has to deal with is drastically reduced. Moreover, different physicians even within the same specialty have different patterns in how they treat patients. This means that reducing the number of physicians that a particular pharmacist or nurse has to interact with reduces the complexity that they have to address. According to a colleague of mine, in Japan pharmacists tend to work with only a few local physicians. Such an approach (with the same set of possible medications) would lead naturally to errors being far less likely, since the possibilities for each pharmacist are drastically reduced.

Creating such specialty teams is not always practical. Still there are other ways to simplify the task of a pharmacist. The basic idea is to separate, as much as possible, the tasks into well-defined and distinct subsets, increasing the effective distance between the tasks even if they have to be performed by the same individual. One way to do this is to separate the pharmacy itself into different areas corresponding to physician specialties. If the specialty of the prescribing physician were marked on the prescription pad (colored specialty pads might be used), the pharmacist would go to the part of the pharmacy with medications for that specialty.

The idea of using teams is also relevant when the standard form of specialization is not sufficient to deal with very high complexity tasks. Physician teams with collective decision-making and action are able to address much more complex tasks than individual physician specialists. Separating a single task to a number of specialists allows them together to perform tasks that can have the sum of the number of possibilities that each one of them faces. Setting up a physician team allows them together to perform tasks that have as many possibilities as the product of the number of possibilities that each one can address. This is a tremendously greater complexity. This is an ideal. It assumes that they all work in a mutually complementary way. Even without perfect complementarity, with proper training, they can work on tasks that are substantially more complex than individuals working separately.

How much specialization and collective action is appropriate—and for what specialties? There's no one answer to this question. Indeed, every hospital or clinic faces a unique flow of patients. The problem of specialization is also linked to the number of cases of a particular type that a medical system sees. Common cases should be treated in a streamlined way, at the other extreme, very rare cases should be treated as exceptions. The effort on a per case basis should increase gradually with how rare the type of case is. The formation of teams, therefore, combines considerations of efficiency and complexity. Specialization should be established so as to best fit the complexity of the medical care required.

What kind of success can we aim for?

Obviously a hospital cannot implement radical structural and organizational changes at once, and the impact of changes on costs is crucial. How to gradually transform an organization into the most effective structure for the complexity of its tasks is the ultimate point of this series of papers and will be discussed in detail in the third paper of this series. The problem associated with prescription errors, however, may be more directly addressed as it is likely to be first and foremost a communication channel problem, due to the convergence of multiple channels from different physicians to the pharmacist. As such, there are very small and easily implemented changes that can be made to reduce this kind of error. What kind of success can we hope for with these changes?

In 1999, the Institute of Medicine's report urged a national goal of reducing medical errors by 50 percent over the next five years. A lot of doctors and health care officials, even those who thought the IOM's medical error statistics were overestimated, thought that this target was overly ambitious. Indeed, today the reduction of medical errors seems far away. How does it sound to you? Well, let's roughly estimate what kind of reduction in communication channel errors we'd get, just using the simple technique of adding redundancy to recording practices.

Studies indicate that a patient admitted to a hospital has a 5-10% chance of being the victim of some kind of life threatening medical error. One simple change can reduce this quite drastically. If there are on average about 10 procedures carried out on each patient during this stay in the hospital, then there is roughly a 1% rate of error for a particular act. (This assumes that the errors are independent. [Aside: The way to calculate this is to ask what is the probability of no error occurring which is $0.90=(1-.01)^{10}$]) Say we introduce some redundancy into the system by having the doctor produce two copies of the prescription, which are checked against each other before administering the drug to the patient. We assume that the errors in each copy are independent, so that each of them has the same individual error rate. By adding this one act (double checking the prescription), the error rate will be squared, and you'll end up with only a 0.01% chance of error for an individual procedure—and therefore a 0.1% chance of error for a particular patient. So with this one small procedural change, we have reduced a patient's chances of being subject to an error by 99%!

This is a simple calculation that does not take into account a variety of factors. Some of these factors would reduce the eventual error rate still further, others would increase it. For example, if the average number of procedures performed on each patient in the hospital is higher than 10, the reduction in error rates would be even greater. If errors are not independent because the people are too tired to write or read effectively then the error rate will be higher. Still the message should be clear: it is possible for a very simple change that addresses the actual problem to have a major impact on error rates, even making the error rate so small as to be unnoticeable. So we've just learned something extremely important: the amount of redundancy that you have to introduce into a system in order to reduce errors to the point of undetectability is not large. You don't have to implement a whole slew of radical changes in procedure in order to dramatically reduce error rates.

Government agencies and independent health safety organizations have very diligently proposed lists of recommendations for changes to address medical errors and some hospitals have responded by spending lots of money and manpower on implementing many of them in a "coordinated attack" on medical errors. Other hospitals have become overwhelmed by the problem of implementing these recommendations. A special emphasis has been placed on technology and automation. It is important to realize that different recommendations will be appropriate for different hospitals, though some changes are likely to be useful for most hospitals. Differences in patient population, physician expertise and nursing programs, may result in a different space of possibilities for the same task at different hospitals—and the methods of error reduction that will be most effective will vary accordingly.

Though the choice of error reduction methods might be bewildering, the effects of appropriate ones, by this calculation, are exceptionally simple. All an individual hospital has to do, then, is pick one or two or maybe even three methods that address the particular communication channel problems they are facing. These changes should bring about a rapid adjustment and near undetectability in a short period of time. The moral of the story is that individual hospitals can try to implement reasonable changes and expect that they will lead to substantial and observable results. At the level of the individual hospital, 50% over five years is absolutely too modest a goal for reducing medical errors.

In the wake of the 1999 IOM report, the reaction of the health care and regulatory community was to focus its efforts on effecting change, originally through centralized action. This task was daunting, and ultimately unlikely to be successful, precisely because its goal is to produce recommendations and procedural changes that would bring all hospitals into line through stricter standards: standardized treatment policies and protocols, and technological devices that would reduce reliance on handwriting and memory. It's not that reducing errors by 50% over five years is not possible, but rather that because the medical system is a highly complex system effective change should arise from within, guided by an understanding of what goals can be reached and what approaches should be tried. Externally imposed standards and regulations will not result in a versatile system that can deal with the complexity of medical needs.

This places the onus on individual hospitals to test new ideas and evaluate them quickly. This fact has increasingly been realized. In March 2001, the IOM released a new report arguing that the health care system had to be reinvented, through a "sweeping redesign" of the entire system—not via the imposition of a "blueprint" for care delivery systems, but through the creative implementation of new simple principles of care. Eager to foster all promising routes for innovation, the IOM has now refrained from specifying proper procedures. It remains to be seen how the promise of the new decentralized approach plays out, but it is a step towards understanding how changes in the complex health care system can be effected. Encouraging local experiments will allow innovative new approaches in health care to be discovered.

Conclusion

Although in this paper, we've focused on errors concerning drug prescription and delivery, the basic insights are profoundly important for other kinds of errors in treatment, misuse or failure of equipment, and incorrect diagnoses: medical errors have to do with complexity. To dramatically reduce the incidence of errors, one must identify where the complexity arises and create a system that has adequately complex capabilities.

Using the notions of complexity and scale, we can get a sense of what a successful medical organization would look like. In a successful organization, the convergence of messages from different types of individuals to one person is limited. Each communication channel is sufficient for the information flow. Unnecessary steps have been eliminated / automated where possible. Standardization reduces the complexity of tasks when it doesn't limit effectiveness. When complexity is unavoidably present, redundancies exist in the system to catch errors. The distribution of complexity across multiple individuals makes it possible for complex tasks to be performed effectively. More specifically, high complexity care is provided by teams with specialization of members of the team, as well as specialization of teams. More self-contained teams provide more individualized medical care from intake through diagnosis, treatment, release and follow up. In this way the traditional reduction of complexity by specialization at the level of the diagnosing specialist physician is maintained throughout the rest of the patient's care involving nurses, technicians, and pharmacists.

In this paper we have focused on how we can think about and design medical services and the teams that provide them. It is quite hard, however, to understand the full complexity of these systems. Rather than designing them, the main role of management and policy makers should be to create an environment in which the systems create themselves. The traditional way to do this is through economic competition. For the healthcare system a different approach is needed, and this will be described in the third paper in this series.

References

[1] World Health Report 2000, World Health Organization, 2001

[2] Institute of Medicine. (2001). Crossing the quality chasm: A new health system for the twenty-first century. Washington, DC: National Academy Press.

[3] Bar-Yam, Y., Dynamics of Complex Systems, (Perseus Press, Cambridge, 1997)

[4] Bar-Yam, Y., General Features of Complex Systems, in Encyclopedia of Life Support Systems (EOLSS), developed under the Auspices of UNESCO, (EOLSS Publishers, Oxford ,UK, 2004)

[5] Bar-Yam, Y., Multiscale Variety in Complex Systems, Complexity, 9, 37 (2004)

[6] Goldstein, A., Overdose Kills Girl at Children's Hospital, Washington Post, April 20, 2001

[7] M. Smith and C. Feied, http://www.necsi.org/guide/examples/er.html

[8] A more targeted approach would focus on the specific drugs that might be confused with each other, posting a short list of the most common name-confusion culprits in pharmacies. The pharmacist would know to watch out for these drugs and when faced with a prescription for one of them, would check the clarity of the writing, or that the prescribing physician is of the right specialty. This would be an "exception handling" system: a system that deals with the cases that are likely to be problematic in a special way. However, exception handling systems place an additional large burden on the person responsible for checking for the exceptions. To be effective, every prescription must be checked against the list. If the list of exceptions includes only one or two cases, this might be reasonable. Otherwise it becomes a significantly complex task in and of itself. If the pharmacist’s existing task is already highly complex, this approach would not be a good one. Moreover, if a prescription was not written clearly, they might have to go back to the physician to ask what was desired, and such inquiries would involve additional delays and increase the burden on the physician.

[9] Horn, S. D., P. D. Sharkey, D. M. Tracy, C. E. Horn, B. James, and F. Goodwin, “Intended and Unintended Consequences of HMO Cost-Containment Strategies: Results From the Managed Care Outcomes Project,” Am J Man Care, Vol. 2, 1996, pp. 253–264; Horn, S. D., P. D. Sharkey, and C. Phillips-Harris, “Formulary Limitations and the Elderly: Results From the Managed Care Outcomes Project,” Am J Man Care, Vol. 4, No. 8, 1998, pp. 1104–1113.