Special to Government Technology Magazine

Automation Fosters Health Care Fraud

Dr. Malcolm K. Sparrow

[Last Revised: December 7th 1996]

With the election over, federal and state policymakers must now face the daunting and urgent task of controlling costs within our major public health care programs. Nearly 62 million people are projected to receive services paid for by Medicare or Medicaid in fiscal year 1996. For these two programs combined, HCFA now pays out roughly $1 billion every working day of the year ($248.9 billion in 1995)—roughly the same volume of money that flows over VISA's entire global network on a busy day.

Searching for tangible cost savings, the health care industry has turned to automation, which eliminates all the paper and most of the people previously involved in claims processing. By the year 2000, virtually all Medicare claims will be submitted, handled, and paid, electronically—the vast majority passing through "auto-adjudication" without any form of human intervention at all.

As we approach the turn of the millennium, public finds are disbursed more and more through massive, fully automated payment systems. This trend affects a broad range of programs aside from health care; including social security, welfare, student loans, tax refunds, and the distribution of funds for emergency aid. Hence the importance of making sure we understand, and respond appropriately to, the effects such systems have on fraud and on fraud control.

Guiding principles: customer-service and efficiency

Automated payment systems are all designed with the important values of efficiency, customer service, and timeliness in mind. And therein lies their appeal for fraud perpetrators, who find such systems very attractive as targets. These systems pay quickly. They *assume* the information presented in the claim is true. And there is little or no risk of human scrutiny. Should a claim be rejected, the system will provide computer-generated explanatory notices, which help make the inner workings of the system transparent to the fraud perpetrator.

Within the health care industry, the art of fraud control is very poorly understood. That is nobody's fault in particular; rather it results from the fact that fraud control, in such an environment, is an extremely complex task, and one that has been sorely underrated. Of course, fraud control would be easy if we could dispense with the competing values of efficiency, timeliness, and service; then we could simply slow everything down, and check everything. But the contemporary challenge is to design effective fraud control strategies that do not sacrifice the values of efficiency, speed, and customer service, and which make sense in a highly automated environment.

Fraud control was always difficult, involving intellectual challenges that have received little attention either from policymakers or from academia. There are no manuals on how to do it. And officials responsible for it find themselves at once isolated (from all their colleagues who serve the ends of efficiency) and besieged (by the apparently inexhaustible supply of thieves trying to game the system).

How much is lost to fraud within our health care system remains unknown. In 1992 the GAO estimated fraud losses at 10%, which would make it a $100 billion problem nationally. But the GAO estimate has no basis in fact, because nobody systematically measures the problem.

Whenever the levels of fraud within high-volume payment systems are systematically measured, the results are usually quite shocking (probably one of the major reasons why officials within the health care industry seem reluctant to commit themselves, or their programs, to rigorous measurement). The IRS, by 1994, suspected they had a major problem with fraud within the electronic filing program. To find out just how bad things had become, the IRS, during the 1994 filing season, picked a random sample of tax refund requests (focusing on refunds based upon the Earned Income Tax Credit, which is where the problem lay), and dispatched criminal investigators to test the validity of each claim. The results of the survey showed 38.8% of the EITC claims were either inflated or entirely unmerited, with 26.1% of the total EITC budget going into the wrong hands. Even with the most conservative definitions of fraud, 19% of the EITC claims were classified as outright fraud. The news was so bad that, for a moment, IRS officials were tempted to discredit their own study. But they resisted the temptation, faced up to the magnitude of the problem, and demanded the resources from Congress necessary to stem the flow.

A statistically valid random measurement program was also used recently by the auto-insurance industry. Insurers reviewed a random sample of 15,000 accident claims from 1992 and found evidence of fraud in 32% of the claims.

Is it possible that the fraud losses in our major health care programs run at similar levels—30% to 40%—significantly higher than prevailing industry estimates? If so, getting to grips with fraud represents a major opportunity to save program dollars without *any* loss of benefits for those that genuinely need them.

When you examine the policies, systems and assumptions that underlie the industry's current approach to fraud control, it becomes quite clear that fraud remains largely uncontrolled. There are virtually no meaningful defenses in place, and therefore every reason to believe that fraud losses run at a higher rate than anyone knows.

Claims payment systems still operate basically on trust and assume that the information content of each claim submitted is true. Accepting that information, automated edits and audits then check to see that the procedure makes sense given the diagnosis, that the pricing is within limits, and that the patient, procedure and the provider are all within the limits of policy coverage. In other words, fraud perpetrators are free to lie, provided they bill *correctly*. Hence

the growth of what some investigators call "perfect paper schemes", or the use of "canned claims", where perpetrators explore the claims payment system and find perfect combinations of diagnosis, procedure, and price that guarantee payment without any kind of human review. Of course, the diagnosis may be fictitious, the procedure might never have occurred, and the patient may know nothing about the claim submitted in their name. But those factors are irrelevant, because nothing in the payment system routinely checks them.

The industry broadly relies upon a set of controls which detect billing errors and medical unorthodoxy, but which generally do not detect fraud. Provided fraud schemes do not produce anomalous billings or patterns (and, of course, the more sophisticated ones do not) then they will remain invisible, not only at the time of payment, but forever. Most of what these systems lose to fraud, nobody but the perpetrator ever knows about. And even the most professional and diligent Special Investigative Unit cannot really get to the heart of the fraud problem if detection methods and referral systems only produce for them the merest trickle of cases when compared to the underlying size of the problem.

Effects of automation

When the general public think about computer crimes, they often imagine hackers breaking into mainframe systems, taking control, and maybe manipulating payments to their own advantage. In fact such crimes are comparatively rare, and are never the major threat in implementing electronic payment systems. The major worry—the one which should make us all rather uncomfortable—is that the system will work perfectly, fast and efficiently, time after time, claim after claim; but with incoming claims which are themselves false.

The surprising motto, for hi-tech fraud, is that *fraud works best when processing systems work perfectly*. The fraud perpetrator wants to know that if a particular claim was "auto-adjudicated" and paid today, then 10,000 similar claims submitted next week will each be processed exactly the same way. Fraud perpetrators love "quality controls" because they help guarantee procedural uniformity, making the payment system perfectly predictable. Quality controls, and other procedural audits, do nothing to test the veracity of the claims themselves.

The serious risks in the electronic environment relate to detection, and the loss of detection opportunities. In particular:

(a) **Absence of common sense in claims review**: the absence of human involvement in reviewing claims means the absence of applied common sense. In paper-based processes, claims that were patently absurd, visibly peculiar, or strikingly similar to others recently seen, would be set aside for review, and human beings had a variety of opportunities to notice the unusual, to spot patterns (even patterns they were not looking for), and to become suspicious. By contrast, claims submitted electronically will be filtered out for review only if the particular absurdity had been predicted in advance and built into the automated checks.

(b) **Computer-generated schemes**: technically competent fraud perpetrators will be able to use computers to generate and dispatch thousands upon thousands of claims, each one designed so as not to attract attention or review, and with the activity spread across hundreds or

thousands of patient accounts so as to avoid detection. In case some think this is crime fiction, it is worth pointing out that such schemes have been around for a decade or more. In November 1988 a New York court convicted the director of one medical center for falsely billing Medicaid for close to 400,000 phantom patient visits between 1980 and 1987. The center's computers were programmed to generate the phony claims and backup medical charts for roughly 12,000 fictitious visits per month.

How is the industry responding to these new threats, and how is it possible to protect against these enormously expensive fraud threats without slowing the payment systems down so much they grind to a halt? The health care industry has given this question some serious consideration, and appears to be headed towards a rather dangerous answer.

"Automated Prevention"

The modern trend is to invest in automated controls rather than additional staff: to put faith in machines, not people. The prevailing beliefs about how to make these systems safe combine two ideas. First, that machines can do a lot more monitoring, and more cheaply, than people can. Second, that prevention is better than a cure.

When you put these two together you end up with a vision of a system which identifies fraudulent claims electronically and kicks them out up front. Call it "automated prevention". The idea is that totally electronic claims-processing systems can be protected from fraud by implementing comprehensive batteries of up-front edits and audits that will keep fraudulent claims out of the system altogether.

This model, unfortunately, is fatally flawed. It assumes that a particular set of automated controls (edits and audits), once implemented, can provide adequate protection against fraud. That assumption underestimates the opposition and ignore the fundamental nature of the fraud control business. In particular:

(a) **fraud control is dynamic, not static**. Given any set of static controls, perpetrators will test them, and learn to circumnavigate them, usually in less than a week.

(b) **transaction-level monitoring is inadequate**. It is not usually possible to distinguish a fraudulent claim from a legitimate one on the basis of its information content alone. Sometimes a broader context may suggest fraud (i.e. a pattern of claims), but the only reliable determination involves external validation--checking with the patients, or with referring physicians.

(c) **lack of useful intelligence**. Automated defenses that rely mainly on "auto-rejection" provide the perpetrator with complete information about what the detection systems can and cannot see. At the same time, they provide little or no opportunity for anyone inside the organization to gather any intelligence about what the fraud perpetrators are doing. Without a human fraud-control operation to do the analysis, only one side in this game is gathering any useful intelligence.

(d) **Auto-rejection is a weak fraud control tactic**. It leaves the perpetrator unscathed, free to try something different tomorrow. Defending systems have to deliver a sting of some kind that will make perpetrators wary of attacking the same target again.

(e) **Lack of identified fraud-control responsibility**. The pervasive vision for fraud control provides little or no place for a human fraud-control team as such. In the absence of such a team, who will be responsible for gathering information about emerging fraud threats and coordinating effective responses? If, in the future, fully automatic fraud-prevention systems really do have no place for human strategists, then the advent of electronic processing will cement in place two of the major failings of fraud control systems today: no one is in charge, and no-one is responsible for fraud control.

"Automated prevention" will fail principally because it will be thoroughly predictable. Perfect predictability makes the target static, transparent, and easy to attack. Effective fraud controls require *unpredictability*, an element of mystery, and have to put the fraud perpetrator at some substantial risk.

The remaining challenge

So the challenge of designing effective fraud control strategies without sacrificing efficiency, speed, and customer service, remains. This is not a simple challenge. The strategy needed cannot be predominantly reactive (as at present), nor completely preventative (as per "automated prevention"). What's needed is a proactive, intelligence-gathering strategy that emphasizes the importance of spotting emerging fraud problems early, and then brings a coordinated, cross-functional approach to eliminating them before much damage is done.

This is not the place to explore such a strategy in detail, but some key lessons for the automated environment are already clear. Here are a few of the core requirements for fraud control in highly automated payment systems:

(a) **Risk of Random Review**. There *has* to be an element of unpredictability. Every claim submitted must suffer at least some small risk of random selection for human review and verification. The risk might be as low as 1%, but it must never be zero.

(b) **External Validation**. The review process must include external verification, and be sufficiently rigorous to uncover fraud if present. At a minimum, verification should normally involve contact with the patient or relatives (if the claim is submitted by a provider), or with the provider (if submitted by the patient).

(c) **Human fraud-control team**. A human fraud-control team should be allowed to operate up front, and pre-payment. They should have day-by-day control over claims-selection criteria so they can arrange for suspension and review of any categories of claims they want to see. External validation of claims should be a normal part of their operation. The fraud control team should never be confused with other systems aimed at billing correctness or medical orthodoxy—the task is quite different.

(d) **Broad range of analytic tools**. The fraud control team should have all the technical tools they need to be able to launch ad-hoc queries within the claims databases. They need to be able to interrogate paid claims, rejected claims, and claims pending payment.

These are just a few of the minimal requirements. Yet this kind of approach differs markedly from the health care industry's current practices. Most insurers, public and private, do no systematic measurement of the fraud problem, which therefore remains largely invisible, and continue to underinvest in controls by a factor of twenty or more. Most insurers fail to designate responsibility for control, and many equate control with investigation, allowing their investigative units to remain bogged down in a reactive, case-making mode, their workload fed to them by ineffective detection and referral systems. The "automated prevention" vision threatens to eliminate human beings from the fraud-control operation almost entirely, and may decimate investigative and enforcement capabilities.

So there remains much work to be done in designing and implementing control strategies that match the magnitude and seriousness of the electronic fraud threat. A major piece of that work involves figuring out the right relationship between people and technology within a fraud control operation. Another piece involves persuading policymakers to balance their current preoccupation with administrative efficiencies by paying more attention to the issues of payment safeguards, prudence, and caution.

Dr. Sparrow, formerly a Detective Chief Inspector with the British Police Service, now teaches at Harvard's John F. Kennedy School of Government. He is the author of "License to Steal: Why Fraud Plagues America's Health Care System", Westview Press, 1996. (Available from Harper Collins on 1-800-242-7737. Paperback $17.95)