Regulation of Bad Things that Almost Never Happen but Could: HIPAA and the Individual Insurance Market

Mark V. Pauly

Since the collapse of the Clinton plan for large-scale health reform, Congress has approached the medical care sector of the economy very gingerly. One apparent lesson of the reform (though an obvious lesson in life) is that people are much less eager for improvements they have to pay for than for those that are supposed to come at zero cost. Once it became clear that the Clinton plan would have had serious financial and distributive implications, the electorate seemed in no mood to even risk the chance of large-scale taxes and transfers.

But the temptation to court votes by "doing something" about the medical care sector, and the obvious and real problems of the sector itself, did, after a several-year hiatus, gradually bring politicians back into full concern mode. First, a few small laws were passed about maternity and mental health benefit design, then came the Health Insurance Portability and Accountability Act (HIPAA), and now there is a full rhetorical court press concerning the Patients' Bill of Rights (PBOR).

Politicians' desire to avoid spending tax collections that would be used for other purposes (including tax cuts) has obviously shaped the form of these interventions. It is the shape of health regulatory intervention that, after some general remarks about its character, I want to consider in the context of some (though by no means all) aspects of HIPAA. My view is that much of that law (as with much of PBOR, possibly even including the liability issue) largely represents regulation that forbids practices that rarely happen. While some of these practices may always be harmful, and others might selectively do harm, my point is that the arguments for these regulations are

Cato Journal, Vol. 22, No. 1 (Spring/Summer 2002).

Mark V. Pauly is Bendheim Professor of Health Care Systems at the Wharton School, University of Pennsylvania.

close to insignificant because of the rarity of the practices they are intended to regulate.

If this is so, one might then ask, on the one hand, why such laws are passed, and, on the other, what is so harmful if they are? Perhaps the only real consequence of note is the consumption of legislative and political energy and resources: a non-trivial deadweight loss if that political time and energy could have been devoted to other uses valued by citizens, but perhaps no more (and even less) of a waste than if it had been used for other political purposes.

My conclusion is less benign. I think that there are potential negative effects for such regulations that ought to be considered, in prospect or retrospect. At a minimum, something might be able to be done to reduce the damage.

My argument here will largely be one of political economy, in which I do have some general expertise. However, I willingly concede that I am an expert neither in the details of the passage of the HIPAA legislation (I have no insight into "what Congress was really trying to do"), nor in the specifics of what effects the legislation has had. Few others, I suspect, know the former, and no one knows the latter. Apart from a few preliminary reports on some aspects of implementation of the law, there has been, to my knowledge, no definitive large-scale evaluation of its effects. Speculation and anecdote, and even judgments of industry insiders, do not constitute reliable evidence.

The Normative Economics of Regulation

Real-world markets are never as perfect as those diagrammed in economics textbooks, and so in theory there can always be some government intervention that could improve matters. However, both common sense and deep theorizing about political philosophy tell us not to hope for a government bent on steadying every shake and shiver in the invisible hand. Usually, "pretty good" is good enough for markets to appropriately remain unregulated.

The main deviation from perfection in medical markets is not usually on the supply side; there is little natural monopoly and, except for pharmaceuticals, few unexploited economies of scale.[1] Instead, the most serious defect, in theory and in the back of all policymakers' minds, is imperfect information: the typical consumer is acutely conscious that he doesn't really know what he is doing when he goes to

[1]Because the large, upfront research and development costs for drugs are fixed costs, the average cost of pharmaceutical products will fall as volume grows.

buy insurance or physician services. Indeed, were he as omniscient as the textbook consumer of widgets, he could frequently avoid buying either.

The case for government intervention usually requires, however, both that consumers not know and that government know better. To be specific: the strongest case for government intervention is to forbid some action that (virtually) no informed consumer who was rational and unforced would choose to do. If "you'd have to be crazy" to buy a particular product if you knew the score, then government, by banning the product, can do consumers the great favor of avoiding the need to protect against or investigate such products, as well as the occasional mishap that seeps through. Exploding cans of shaving cream, collapsing ladders, and de-laminating tires all fit into this category.

A trickier call is the case of products that some (if fully informed) consumers like while others (if fully informed) hate. The latter group would gain from a ban on such products, while the former group would suffer. A straightforward and utilitarian approach would total the gains and losses and then take the action that maximizes the net. Generally speaking, however, this second case is a better candidate for a full-court press on consumer information (indicate in bold print which dishes are extra spicy) rather than those sorts of Benthamite calculations. But if the many are really many and the few really few, and information hard to provide, there might be a case. So the correct condition might be labeled "almost everyone would have to be crazy."

Necessary and Sufficient

The previous discussion is correct and helpful as far as it goes, and it is a good guide to cases where product quality regulation is not a good idea (e.g., when only a few people are unfamiliar with the dangers of chain saws or butane). However, the "almost everyone" standard would be both necessary and sufficient for intervention only if we could count on government to do the right thing well (and only the right thing). But there are reasons why that may not be a reasonable expectation. I will deal with three of these. (There are more.)

1. The common requirement to certify individuals, organizations, and entities as in conformance with the regulations can impose substantial administrative costs on the certifiers and certifiees.
2. The existence of rules and regulations often calls forth its counterpart: rent-seeking behavior engendered by the rules, usually (though not always) via the legal system. If there are rules, then

the threat to bring someone to court for violations of those rules, the extension of the rules to other similar behavior which might fall under them, and the need to check the coordination of federal rules with other state or local rules all provide fertile ground for resources to be deployed.

3. Finally, regulations that do an imperfect job of defining behavior are, at a minimum, confusing, and at worst may lead to a degradation of behavior down to the minimum standard required by law. Just as there must be far too many inefficient meetings at Grand Central Station because it serves as a fixed point, a poorly drafted or incomplete regulation may play the same role.

In what follows, I will apply these three concepts to two features of HIPAA, one central and well-known and the other obscure but deserving of notice: preexisting conditions exclusions in employment-based group insurance and guaranteed renewability provisions in individual insurance.

Preexisting Conditions Exclusions, Portability, and HIPAA

One of the primary goals of HIPAA was to allow workers who had insurance coverage on one job to move to another job without a break in coverage for continuing conditions. Both the loss of coverage for those who did move, and the "job lock" allegedly experienced by those who could move but did not, were thought to be defects of the traditional employment-based system.

From the "almost everybody" standard, preexisting conditions are a legitimate concern. No one would choose to have insurance coverage that disappeared for an ongoing condition just because they changed jobs. Ideally, the way to handle this problem would be for insurance to be fully portable across jobs. If we cannot avoid the employment-based setting, the preferred method would be for the current employer's insurance (or continuation insurance paid for out of compensation for the initial job) to continue to cover the condition until the new employer's coverage kicked in. (I cannot avoid noting at this point that this whole mess is itself caused by the system of tax-subsidized employment-group insurance. In contrast to health insurance, the worker in the midst of settling claims on his auto insurance does not lose collision coverage temporarily when he switches jobs—because auto insurance is independent of the job.)

What HIPAA did was something different than what would have been efficient—it limited the duration of the exclusion to no more than a year, and required the new employer to cover the "runout" of

the ongoing condition. The duration and applicability of exclusions were more limited for those workers who did not have sufficient prior coverage.

Was this sensible policy? I do not want to rehash the HIPAA debate at length, but it is worth noting that (a) many employers offering coverage had no or brief waiting periods, (b) few workers have expensive ongoing conditions in their families, (c) even fewer of the workers that do are those who change jobs, and (d) COBRA rules would make it rational for workers to try to continue their coverage if the uncovered expected expense exceeded the COBRA premiums (a strategy that many workers can play ex post facto for a limited period of time—60 days—after they leave previous job).

While we have no definitive estimates, I would expect that less than one percent of all workers in a given year would obtain benefit from the HIPAA rule. If we start with the (overly pessimistic) GAO estimate that about eight percent of all workers and dependents "could be affected" because they change jobs, we must then subtract all those who took new jobs without exclusions, without insurance, or with waiting periods, and delete those who could switch to spouse's coverage. This would reduce the percentage affected to two or three percent. But of that number, depending on one's definition, only a very small fraction would experience a non-postponable need for a medical service for the preexisting condition.

The most serious cost of this provision so far has been the expense to employers to certify, for any worker who changes jobs, regardless of health status or rules at the new employer, that the person had x months of qualified coverage. After I visited the Economics Department at Stanford for six months and returned to Penn with full benefits, Stanford found me and sent me such a certificate for my wife and myself. (They could not find me to mail my W-2 form.)

We do not as yet have evidence on whether this provision will generate serious litigation. Probably the most gaping hole in the promise of portable coverage is that the law has zero effect if the new employer does not offer insurance and the employee does not seek expensive individual coverage. If the employee does, there is a requirement to provide some type of individual conversion coverage but, in most states, no rule about what the premium could be. Probably equally important, the law is ineffective if the new job provides insurance, but there is a waiting period for any coverage. Gabel et al. (2001) recently showed that waiting periods of four months or more are common, especially among smaller firms.

So has this provision greatly improved the portability of coverage and substantially reduced job lock? No one can be sure, to any

appreciable extent. There was not a (proportionately) large problem to begin with, and any binding regulation can be avoided by dropping coverage or by instituting waiting periods. Of course, even if the law affected less than one percent, because this is a big country, that would still represent hundreds of thousands of people. But because this is a big country, the cost of certifying compliance with the law would affect millions more.

Guaranteed Renewability in Individual Health Insurance

For people who do not choose to work at firms that offer health insurance coverage, insurance is obtainable in the individual market. In contrast to the group market after HIPAA, individual insurers in most states are permitted to charge premiums related to risk. They do tend to charge more to older persons and less to young persons, and people who report chronic conditions on their application may be charged more or excluded from coverage for their conditions for a period of time. Both of these latter strategies themselves are relatively rare (Pauly and Herring 2000), for reasons that I will now discuss.

A common feature in individual health insurance (and life insurance) that consumers can obtain readily is guaranteed renewability (GR). Perhaps 80 percent of health insurance policies before the passage of HIPAA were guaranteed renewable. A GR provision protects an insured individual against increases in premiums related to changes in that individual's health risk. Insurers promise that they will only raise premiums for all individuals in a rating class, not selectively for specific individuals based on their health status or anything else.

While GR does not provide bulletproof protection—premiums can increase, insurers can exit a market entirely, and a rating class can become overloaded with high risks—it still helps a great deal. It probably explains why many high-risk people are able to buy individual health insurance at relatively low premiums (Pauly and Herring 2000).

Despite the fact that a very large fraction of individual plans previously contained explicit guaranteed renewability provisions, and probably more followed the practice of not checking health status at renewal, HIPAA included a requirement to make the feature obligatory for literally all individual insurance. This provision is problematic for two reasons.

First, there are occasionally reasons why "non-crazy" buyers might want a plan without regard to renewability, for example, when they buy insurance for "bridge" coverage between jobs. That insurance would carry lower premiums if it did not guarantee renewability, but

it could not continue to be cheap if people could renew at uniform rates. However, insurers may be able to avoid this problem by having rating classes that terminate in 12 or 18 months.

The second concern is subtler. The formal HIPAA requirements for guaranteed renewability are seriously incomplete. The idea of guaranteed renewability in insurance theory and insurance texts is as follows:

> At the end of the policy period the insurer (under guaranteed renewability) must renew coverage regardless of the health of the insured. While the premium can increase to reflect the expected experience for the individual's rating class, the premium cannot be increased on an individual basis to reflect possible deterioration in the individual's (or covered dependents') health [Harrington and Niehaus 1999: 460].

The HIPAA law defines guaranteed renewability in the same way as the first sentence in this definition. But it totally omits the second sentence. As a result, it allows for an insurer to guarantee renewability for high risks but propose to charge them an infinite premium, effectively vitiating any protection. While other state laws may limit risk rating, it still appears to be the case that some states' legislation to accommodate HIPAA permits this loophole.

Why, for someone skeptical about the need for regulation in the first place, would the laxity be a matter of concern? After all, HIPAA does not prevent any insurer from defining guaranteed renewability in its contract in the way the textbook does, and the way most insurers behaved before the law. My concern arises from an expectation that government rules, by setting a "clear and uniform standard," will tend to push contracts toward that standard, *even when the standard is lower than what most buyers in the market would demand.* Before there were rules, buyers (or their brokers) would naturally be expected to explain and explore carefully all aspects of contract language; with the rule, there can be a false sense of security. (A recent parallel is the security about medical quality that most people feel from physician licensure contrasted with a market allegedly fraught with errors and defects in practice.)

My concern might be viewed as excessive pessimism (or even paranoia), were there not a superior remedy available. Rather than pass a vacuous regulatory requirement, it would be better for regulation to require sellers to disseminate easily noticeable and understandable information about such a key policy feature, and then let buyers decide what is "not crazy" for them. False security is worse than no security at all. Of course, had there been a serious problem with the

way guaranteed renewability for individual insurance was working, and had the HIPAA law addressed that problem, there is a chance that the legislation would do more good than harm—but neither of these premises appears to be true. Indeed, although many states regulated guaranteed renewability for small group insurance before HIPAA, there appeared to be very little regulation of this policy provision in the individual market, with only a handful of states even requiring guaranteed renewability in some form or other.

As far as I know, there has never been a study of how guaranteed renewability actually works in the individual market, and even the data I have mentioned are ambiguous and imprecise. Why then, one might ask, should this provision be expected to be common and to work well, especially when it seems to have been uncommon enough to generate state regulatory requirements in the apparently similar small group market? I think answers to these somewhat esoteric (if nevertheless intriguing) questions shed light on the appropriate role of product quality/insurance quality legislation overall.

There are (at least) two reasons why guaranteed renewability emerged naturally in individual markets and not in group markets. One is suggested by Harrington and Niehaus. A rating class begins with all members at roughly the same health level, but over time some individuals or small groups will develop more severe problems than others in the rating class. There is then a strong temptation for the healthy individuals or firms to conspire with other insurers to offer them premiums less than the premiums that would be needed to cover everyone in the class; the low risks are (quite willingly) picked off, and premiums rise uniformly but rapidly for those who are left. Harrington and Niehaus (1999) suggest that employers are likely to be more interested in trolling for favorable rates than individual consumers (who have a life) would be. In this case, inertia is our friend, and protects us against cream skimming in the individual market.

The other reason is that individuals may have a stronger demand for and a lower cost of guaranteed renewability protection than small employers do. For a risk averse individual, the onset of a chronic condition has two adverse financial effects. There will be unusually high medical bills in the year in which the condition is contracted, and there will be a dismal future of above-average insurance premiums in the absence of guaranteed renewability. Conventional insurance shields the individual from the first risk, but not the second. Not only will consumers prefer an insurance contract that guarantees to protect them against the second risk, they will be willing to pay for that protection in a way that prevents or limits the incentives for low risks and insurers to engage in cream skimming.

Here's how. Think of the first period after people have bought guaranteed renewable insurance in a given rating class. The premium they are charged for coverage in that period ought to have two parts. One part, as usual, pays for expenses incurred in that period. But the other part pays for protection for those few people who become high risks. In effect, people "prepay" their above-average future premiums. With this prepayment in place, there is then no need to raise the premium for anyone, including those who remain low risks, in the next period, and therefore no reason for the low risks to drop out of the pool. In short, the problem that arises under what insurers call "durational effects" can be avoided if insurers and buyers plan ahead. Of course, premiums for a rating class that is aging may rise just because expected expenses even for healthy people rise as they get older, but that increase should be totally expected and, more importantly, unlikely to cause the healthy but older members to drop out and seek coverage elsewhere (since any new insurer will notice their age).

Small employers, however, do not have the same kinds of insurance demands as individuals do. For one thing, the insurance is not for them. More importantly, economists believe that employers do not pay for the insurance they offer to employees out of profits but rather out of what would have been their workers' higher wages. A common reason why a small group experiences higher expected expenses this year than last year is that the composition of its workforce changed. Workers may have gotten older on average, and some new hires might have chronic conditions in their families. But that change does not represent so obvious a risk to the employer, who may be able to adjust wages downward to deal with it, and who in any case may be able to ride it out for a year or two until workforce composition changes again. Cutler (1994) did find evidence that some small employers (the ones you would expect) demanded and got protection from their insurers against premium fluctuation over time, but one suspects that many employers attach little value to such protection, which will be hard for insurers to furnish. Insurers can forecast changes in the number of individuals in a given population who will have chronic conditions, but they cannot easily forecast changes for small groups of changing membership.

So now what happens if regulations force insurers to provide a costly policy feature that buyers don't want (given its cost)? The sellers think of ways to slide out from under the obligation, and the buyers don't mind. What seems (based on anecdote) to happen in the small group market is that durational effects sometimes occur and protection is indeed eroded. My hypothesis is that this phenomenon

should be much less common in individual insurance—not because such insurers are nobler, but because individual buyers do not demand the false contracts that many group buyers do.

It is the demand of informed consumers that disciplines the process, in a way that is much more effective and less costly than regulations are likely to be. To make guaranteed renewability provisions work in the individual market, one needs more than just clear contract language. One needs a process of reputation formation or some other method of conveying information to buyers about what sellers do. If an "unscrupulous" seller plans to raise premiums in a rating class, sell a new policy type to the low risks who drop out, and refuse the high risks, but buyers all know about this behavior, no one will purchase from that seller. Much individual insurance is offered by Blue plans, which do appear to have good reputations, and some other sellers in this market are appropriately distrusted by those in the know. The best thing government could do here, as above, is inform, inform, inform.

The main message is that regulations that require sellers to do things that buyers do not really want will inevitably tempt both groups to co-conspire in avoiding these limits. In contrast, contract provisions desired by buyers will be self-enforcing. Of course, the reputational process is not perfect, and this is a big country. But with some assistance in spotlighting behavior intended to evade explicit or implicit policy provisions, regulation to forbid those behaviors may be unnecessary, and counterproductive if it does occur.

Defined Contributions and Other Innovations

Are there future prospects for passing other laws against bad things that almost never happen? In some sense, virtually any innovation falls into this category, since there is no innovation for which a smart person cannot conjecture some bad thing that "might" happen, there being no data on experience to act as a constraint. Moreover, the usual grab bag of problems—principally threats to quality or risk pooling—can always be opened for any innovation in insurance. But the regulatory threat may inappropriately stifle innovation that will either give us better new products or tell us to be satisfied with what we have because it is as good as things are going to get.

Let me offer an example of this, by discussing possible employer movement to defined contribution administration of health benefits and the use of electronic methods for insurance (DC/EI). I have argued elsewhere (Pauly and Given 2001) that these two innovations make the most sense if they are combined, but that even then a future

of large market shares for these innovations is far from a sure thing. If this evaluation is accepted, what policy should be followed for regulating these devices?

From what I have said earlier, I believe it would be wise to go slowly and lightly in this regard, making as much use of existing regulation as possible (since old laws are usually good laws, or at least laws we are used to) for electronic methods, and permitting voluntary agreements between employers and workers to police the form of benefits contracts and administration. (I ignore for the moment the issue of tax treatment of defined contribution plans, an issue that cannot really be ignored in the larger context.)

We strongly suspect that DC/EI will be a boutique benefits product (at least at first, by necessity), so whatever harm it does is bound to be rare. Moreover, the full information model would seem to be better than the tight control model. Finally, it is important for this strategy to be permitted to be tried, so that we can see if workers really want to choose their health insurance themselves, or whether they are willing to delegate that task to their employer's benefits department. Plausible (and even, surprisingly, passionate) arguments can be made either way, but there is no substitute for a neutral market test that can be allowed either to succeed or fail.

Of course, the tax subsidy gets in the way (although the "sue the boss" possibilities from PBOR may artificially stimulate DC/EI). However, as in other industries, it is important that sellers of innovative products be given enough rope, and then we can see where the market goes. Consumers need to be protected from gross harm in the process, but, since employers always have the option of distributing no part of a worker's compensation in the form of health benefits, and since labor markets are usually competitive, it is hard to believe that a DC/EI offering from a job that workers can either take or leave will do much harm, especially with good information.

Conclusion

In the two dimensions I have discussed, HIPAA did embody regulations with the potential for improving efficiency. Its rules about portability of coverage across jobs and guaranteed renewability for individual insurance were targeted at eliminating behaviors which very few knowledgeable buyers would have desired to occur. The legislation in its current embodiment may well have done more good than harm (especially for portability), although we do not yet have and may never have the basis for judging for sure.

My main conclusion, however, is that the probable small-scale na-

ture of the targeted problems suggests that the traditional regulatory character of HIPAA was poorly suited to the task. It may be better to modify the process. That modification should begin by trying to get a better understanding than was available when HIPAA was passed of *why* the targeted behaviors occur. Blaming them on evil desires of employers and insurers (while not always wrong) does not help. Bad behavior usually proceeds either from bad incentives or from bad information (or both). In the case of portability, we know the fundamental bad incentive: the distortive tax subsidy to employment-based group insurance. Both problems—lack of portability and lack of guaranteed renewability—are subject to bad incentives, in that the optimal policy in both cases, a policy of prepayment, is neither permitted by nor fostered by existing state regulations and industry practices.

Finally, in both cases, a bright spotlight on bad behavior—this employer designs benefits to lock you in to a low-quality job; that insurer promises to keep charging you reasonable premiums if you get sick but uses sharp practices to renege on the promises—could be effective. This disclosure strategy seems preferable to designing ponderous regulatory schemes that only a lawyer could (and would) love for behaviors that are rare and relatively inconsequential in the overall scheme of things. After all, the main social problem in health insurance markets is the uninsured, and none of this legislation should have been expected to make a perceptible improvement in that problem. Producing better "access" to insurance through regulation is a mirage. Access requires money. It should be the general taxpayers' money, not that of other unlucky insurance purchasers or even insurance firm owners (as if we thought we could ever really make them pay).

References

Cutler, D. (1994) "Market Failure in Small Group Health Insurance." NBER Working Paper No. 4879 (October).

Gabel, J.; Pickreign, J.; Whitmore, H.; and Shoen, C. (2001) "Embraceable You: How Employers Influence Health Plan Enrollment." *Health Affairs* 20: 196–208.

Harrington, S. E., and Niehaus, G. R. (1999) *Risk Management and Insurance*. Boston: Irwin/McGraw-Hill.

Pauly, M., and Given, R. (2001) "Defined Contributions and E-Health Insurance: Do They Go Together?" Paper presented at Princeton Conference on "The Future of Managed Care," sponsored by the Committee on the Economic Impact of Health System Change (May).

Pauly, M. V., and Herring, B. J. (2000) "An Efficient Employer Strategy for Dealing with Adverse Selection in Multiple-Plan Offerings: An MSA Example." *Journal of Health Economics* 19 (4): 513–28.