Commentary

Health Plan Switching and Attrition Bias in the RAND Health Insurance Experiment

John A. Nyman
University of Minnesota

Abstract One of the most influential "lessons" of the RAND Health Insurance Experiment (HIE) is that cost sharing can reduce hospitalizations by about a quarter, with no effect on health for the average adult. In an earlier paper in this journal, I suggested that a portion of this reduction is due to participants becoming ill and dropping out of the experiment in order to switch to their preexperiment insurance coverage and thus avoid paying the cost-sharing amount. The sixteenfold higher voluntary attrition rate in the cost-sharing arms provides compelling evidence in support of this alternative explanation. Evidence is also provided by the finding in Manning, Duan, Keeler (1993) that the predicted number of hospitalizations among those who dropped out of the coinsurance arms was significantly greater by 34.5 percent than the actual number of hospitalizations, suggesting that participants anticipate hospitalizations and leave the experiment before incurring the cost-sharing payment. Still more evidence is provided by the finding that those (cost-sharing) participants with greater incomes, instead of being *more* likely to be hospitalized, as greater income usually implies, were *less* likely to be hospitalized than poor participants. This finding is consistent with their having better preexperiment insurance coverage than poor participants and therefore being more likely to have an incentive to drop out. This inpatient attrition bias makes it dangerous to rely on this lesson of the HIE, because it is not clear that hospitalizations were actually reduced by one-quarter, much less that if such a reduction actually had occurred, there would be no health consequences.

The RAND Health Insurance Experiment (HIE) concluded that cost sharing can reduce health care spending—specifically hospital use, by about

I would like to thank Bryan Dowd for helpful comments. The views contained within this commentary, plus any errors or oversights, are mine alone.

Journal of Health Politics, Policy and Law, Vol. 33, No. 2, April 2008
DOI 10.1215/03616878-2007-062

one-quarter — with little or no effect on the health of the average adult participant. Reducing hospitalizations by one-quarter would reduce health care expenditures considerably, and because the RAND experiment determined that there were no important health consequences as a result, this "lesson" has been enormously influential on policy makers (Gruber 2006). Over the years, it has been used to justify a variety of proposals to apply cost sharing across the board and without regard to the types of services or the diseases to which it is applied.

In a recent paper in this journal (Nyman 2007), I suggested that much of the difference in health care utilization that is observed between the cost-sharing and free-care arms of the experiment, especially the difference in hospitalizations, could be explained by health plan switching that is evidenced by a rate of voluntary attrition in the cost-sharing arms, which was sixteen times higher than in the free-care arm (Newhouse and the Insurance Experiment Group 1993: 24). This interpretation raises doubts about whether a 25 percent reduction in hospitalizations actually occurred and, furthermore, if such a reduction had occurred, whether it would actually have had no effect on health. Joseph Newhouse and nine of his HIE colleagues have provided a response (Newhouse et al. 2007), and this commentary is my reply.

Speculation

Newhouse et al. (2007) describe my alternative interpretation as mere "speculation," but speculation is what science is about. When confronted with observed data, it is the responsibility of scientists to speculate as to what underlying mechanisms are at work. Speculation about observed phenomena leads to new theories, new empirical tests, and new knowledge. While there is a conventional explanation for the behavior observed in the HIE, others are possible. Alternative explanations should generally be encouraged because they represent the dialogue by which science progresses.

The specific behavior in question is the observed reduction in hospitalizations in the cost-sharing arms. The conventional explanation is that the entire difference is due to consumers purchasing less health care when faced with a higher cost-sharing price. Equally reasonable, however, is the explanation of insurance plan switching, for which there is a sizable supporting literature (Nyman 2007: 773). Participants who become ill and face a large cost-sharing payment drop out of the experiment in order to revert to their more generous preexperiment insurance. The sixteen-

fold higher voluntary attrition rate in the cost-sharing arms is compelling empirical evidence that the plan-switching explanation is not mere speculation.

Counterarguments

Newhouse and the Insurance Experiment Group (1993: 26) concluded that attrition from the experiment did not cause "any appreciable bias." They based this conclusion on two arguments. First, participants in the cost-sharing arms were given participation payments to continue in and complete the experiment. As a result, Newhouse et al. (2007) hold that "there is no a priori reason to expect that financially motivated attrition occurred."

These participation payments may have reduced the incentive to drop out of the cost-sharing plans, but they did not eliminate it. First, large co-payments for an anticipated hospitalization might have presented a cash-flow problem that, even though alternative solutions were available, could have been solved more easily for some simply by leaving the experiment. Indeed, a cash-flow problem caused by the cost-sharing payment may have been exacerbated by the additional nonmedical expenditures and reduced income that are often associated with illness and for which there are fewer alternative solutions. Second, Newhouse et al. (2007) implicitly assume that the value of the future payments from the participation incentive was constant for the individual. However, a person with a health issue is likely not to be as future oriented as one without. Therefore, he or she might have discounted the value of future payments for staying in the experiment—perhaps severely, if the diagnosis was life threatening—compared with the more immediate issue of paying a large co-payment.

The more likely effect of the participation payments would have been to assure that only those who were ill enough to face the prospect of a large medical expenditure (such as a hospital procedure) and the resulting maximal cost-sharing payment would voluntarily exit the experiment. Few who were healthy or relatively so would have left the experiment voluntarily, even if randomized to a cost-sharing arm, because they did not want to forgo these participation payments.

With regard to their second argument, RAND researchers reasoned that, if those who left the experiment were ill, they would have *increased* health care use before leaving. Because the researchers found that "there was no statistically significant difference in the rate of use between those

who did and those who did not complete the Experiment" (Newhouse and the Insurance Experiment Group 1993: 25), they concluded that there was no "appreciable bias." An equally reasonable supposition, however, is that those who left would have *decreased* use in anticipation of receiving the diagnostic tests and hospital procedures outside the experiment. Thus, the combination of increased and decreased use by the leavers might have resulted in the finding of no difference in use, even if attrition bias had occurred.

Newly Posted Study

In response to the publication of my article (Nyman 2007), Newhouse posted on his Harvard University Web site in October 2007 a paper by Manning, Duan, and Keeler, dated April 1993, that draws far different conclusions than those reported by Newhouse and the Insurance Experiment Group (1993: 25–26). In the posted study, the authors regressed outpatient episodes and inpatient admissions on variables representing the insurance plan, health status, and demographic and economic characteristics for all those who completed the experiment. This equation was then used to predict outpatient and inpatient use for those who dropped out of the experiment, based on their characteristics. Expected use was then compared to the actual use by those who dropped out.

The results indicate that those in the coinsurance (or Family Pay) plans who dropped out had 34.5 percent fewer actual admissions than were predicted, and the difference between their actual admissions and their predicted admissions was statistically significant, with a *t*-statistic of –2.61 (Manning, Duan, and Keeler 1993: 12). These results are consistent with the explanation that these participants had anticipated the need for a hospitalization and dropped out before being admitted in order to obtain better insurance coverage outside the experiment. This is clear evidence that attrition had introduced an "appreciable bias" into the published results of inpatient use, contradicting the conclusions in Newhouse and the Insurance Experiment Group (1993: 25, 26).

In their response to Nyman (2007), Newhouse et al. (2007) continue to deny an attrition bias, claiming that I have "no evidence that the difference in attrition is not ignorable (in the statistical meaning of the term), that is, no evidence that it biased the results" and that "we could find no evidence that those who withdrew were importantly different from those who did not." Manning, Duan, and Keeler (1993: 15–16), however, summarize the results of their study in this way:

> The data from the Health Insurance Experiment indicate that individuals who dropped out behaved differently (in ways that measured characteristics and time on the study could not explain) than those who stayed to the end of the study. Thus, past inferences based largely on stayers are biased. For outpatient care, the bias was quite small. But for inpatient care, there was a moderate bias. . . .
>
> The direction of the bias surprised us. We expected that high users on the pay plans would drop out so that they could use their more generous nonexperimental coverage. To avoid this, the HIE made a worst case payment so that no family was financially worse off for being in the study. We expected no bias on the free plan, because its coverage and breadth of benefits was much better than most existing private or public insurance packages. . . . Instead we found that it was high users of inpatient services on the free plan and low users of both inpatient and outpatient care on the family pay plans who dropped out.

Thus, the evidence from this study and the authors' conclusions again directly contradict the assertions of Newhouse et al. (2007). Why Newhouse and colleagues would continue to deny the presence of attrition bias in the face of this evidence is a mystery.[1]

The true level of attrition bias, however, remains elusive. Data on utilization after leaving the experiment are missing, so inferences must be made from the data available. Because illness is a largely random event, the econometric model on which these calculations were based could explain only 14 percent of the variation in hospitalizations (Manning, Duan, and Keeler 1993: 12). Thus, a large amount of variation in the hospitalizations before leaving cannot be explained with the regression equation used by Manning, Duan, and Keeler (1993). As a result, their 34.5 percent differential between the actual and predicted hospitalizations should probably be regarded as a lower bound estimate of the attrition bias.

Both Conventional Demand and Plan Switching

Newhouse et al. (2007) calculate a worst-case scenario that assumes that the entire difference in hospitalizations between the free-care and coin-

1. A paper with the same title and authors as Manning, Duan, and Keeler (1993), but dated 1988, appears in the references of Newhouse and the Insurance Experiment Group (1993: 474), but its results were not referred to in the section of this book on attrition.

surance (Family Pay) plans is attributed to those who dropped out voluntarily and then argue that the resulting additional hospitalizations would be unreasonably large. This is not a position I take in Nyman (2007), because only those with favorable preexperiment insurance would have had the incentive to switch plans and not all had such favorable insurance. Indeed, some did not have any preexperiment insurance at all (Newhouse and the Insurance Experiment Group 1993: 462n8). Therefore, the difference in hospitalizations between the free-care and coinsurance arms probably reflects some combination of the conventional demand explanation and plan switching.

These two responses are evident in table 3.4 from Newhouse and the Insurance Experiment Group (1993: 46), a portion of which is reproduced here as table 1. Table 1 shows that participants in the lower-income group responded to cost sharing as expected for both the probability of any use and the probability of inpatient use. We would normally expect that participants in the medium- and higher-income groups would exhibit a higher probability of health care use, because consumption of medical care generally increases with income. While higher-income participants exhibited a higher probability of consuming any care, they exhibited a lower probability of consuming any inpatient care, consistent with plan switching and dropping out by those who were more likely to have generous preexperiment insurance policies by virtue of their higher incomes. Low-income participants were less likely to have generous preexperiment insurance policies to return to, and so their response would largely reflect the conventional demand explanation (R. Waldmann, personal communication, October 23, 2007). If it was predominantly the poor who actually were forgoing important medical care in response to cost sharing, then it is perhaps no coincidence that the HIE also found that only the poor (who were sick) exhibited a negative health effect from cost sharing (Newhouse and the Health Experiment Group 1993: 208).[2]

New Insurance Theory

In their final paragraph, Newhouse et al. (2007) suggest that the theory presented in the first part of Nyman (2007) is not new but was developed originally by de Meza (1983). De Meza's article, however, was concerned

2. In Nyman (2007), I mistakenly claimed that health-status information had not been collected once the participant had left the experiment, but actually the portion of the health-status data that was self-reported was obtained from 85 percent of living respondents after they left the experiment (Newhouse and the Insurance Experiment Group 1993: 19).

Table 1 Probability of Any Use and Any Inpatient Use by Plan and Income Group

	Income Groups				
	Lower third	Middle third		Higher third	
Plan	Mean	Mean	*t* vs. lower third	Mean	*t* vs. lower third
Probability of Any Use (%)					
Free	82.8	87.4	4.91	90.1	5.90
25%	71.8	80.1	5.45	84.8	6.28
50%	64.7	76.2	4.35	82.3	4.86
95%	61.7	68.9	3.96	73.8	4.64
Individual deductible	65.3	73.9	6.09	79.1	7.09
Probability of Any Inpatient Use (%)					
Free	10.63	10.14	−0.91	10.35	−0.35
25%	10.03	8.44	−2.95	7.97	−2.75
50%	9.08	8.06	−1.78	7.77	−1.66
95%	8.77	7.38	−2.79	7.07	−2.46
Individual deductible	9.26	9.44	+0.31	9.88	+0.68

with insurance that paid beneficiaries "a lump-sum cash payment in the event of illness" (de Meza 1983: 50). De Meza did not consider insurance that paid for the beneficiary's care, the type of insurance actually bought and sold in the United States. This type of insurance effectively reduces the price that consumers face for medical care, and this price reduction was considered the only source of the moral-hazard welfare loss under conventional theory. For this reason, Pauly (1983), in his response to the de Meza (1983) article, wrote that "it is nevertheless true that the relevant theory, empirical evidence and policy analysis for moral hazard in the case of serious illness has not been developed. This is one of the most serious omissions in the current literature" (Pauly 1983: 83).

In Nyman (2003) and other works, I address the issue of how to analyze moral hazard in standard insurance—that is, insurance that pays off by reducing the price of care. I suggest that for most of the costly health care procedures—the care of "serious illness" that Pauly was concerned about—the price decrease is effective only for those who are ill. What *healthy* consumer would purchase a serious procedure like a coronary bypass operation or a liver transplant just because he was insured and the price of care had dropped to zero? Thus, the price reduction is the

vehicle by which income is transferred from those who remain healthy to those who become ill, and as a result, much of the additional health care consumed (that is, much of the moral hazard) is due to an income effect and is efficient. That an income transfer effect was contained in the price reduction of standard insurance was not addressed by de Meza (1983).

The new theory suggests that a portion of moral hazard can still be inefficient (the portion caused by the substitution effect), but another portion—and likely a larger portion (the portion caused by the income transfer)—is efficient. This implies that cost sharing should be applied only to the additional care that is generated by the substitution effect. Cost sharing should not be applied to the portion of the additional care that is generated by the income transfer because it is valuable care that would likely have important health consequences if it were not obtained.

Effect on Health

These implications show why the conventional conclusions of the RAND HIE are potentially so harmful. Not all moral hazard is inefficient. If cost sharing is applied across the board, it does not distinguish between care that is efficient and inefficient, and it discourages both types of care.[3] Discouraging efficient care by applying across-the-board cost sharing could have serious health consequences. The RAND experiment, however, concluded that applying cost sharing across the board reduced care substantially, especially hospital procedures (by about 25 percent), with no important health consequences for the average participant. But if only a portion of this reduction actually occurred, then the conventional conclusions are unreliable and potentially misleading. Therefore, it is important that the inpatient attrition bias in the RAND HIE results be acknowledged.

3. For example, based on RAND Health Insurance Experiment data, Manning and Marquis (1996) concluded that an optimal insurance policy would have a 45 percent coinsurance rate, applied across the board, with no limit on out-of-pocket spending.

References

de Meza, D. 1983. Health Insurance and the Demand For Medical Care. *Journal of Health Economics* 2:47–54.

Gruber, J. 2006. *The Role of Consumer Copayments for Health Care: Lessons from the RAND Health Insurance Experiment and Beyond*. Menlo Park, CA: Kaiser Family Foundation.

Manning, W. G., Jr., N. Duan, and E. B. Keeler. 1993. Attrition Bias in a Randomized Trial of Health Insurance. Unpublished manuscript. www.hcp.med.harvard.edu/files/Attrite_2.pdf (accessed on December 10, 2007).

Manning, W. G., and M. S. Marquis. 1996. Health Insurance: The Tradeoff between Risk Pooling and Moral Hazard. *Journal of Health Economics* 15:513–654.

Newhouse, J. P., R. H. Brook, N. Duan, E. B. Keeler, A. Liebowitz, W. G. Manning, M. S. Marquis, C. N. Morris, C. E. Phelps, and J. E. Rolph. 2007. Attrition in the RAND Health Insurance Experiment: A Response to Nyman. *Journal of Health Politics, Policy and Law* 33:295–308.

Newhouse, J. P., and the Insurance Experiment Group. 1993. *Free For All? Lessons from the RAND Health Insurance Experiment*. Cambridge, MA: Harvard University Press.

Nyman, J. A. 2003. *The Theory of Demand for Health Insurance*. Stanford, CA: Stanford University Press.

———. 2007. American Health Policy: Cracks in the Foundation. *Journal of Health Politics, Policy and Law* 32:759–783.

Pauly, M. 1983. More on Moral Hazard. *Journal of Health Economics* 2:81–85.