

**Attrition in the RAND Health Insurance Experiment:
A Response to Newhouse et al. (2007)**

John A. Nyman
University of Minnesota

October 28, 2007

Introduction

The finding that the imposition of cost-sharing had little or no effect on the health of the average participant as measured by a series of health measures is undoubtedly the most important and influential result of the RAND Health Insurance Experiment for policy. Our confidence in this finding, however, is called into question by the fact that cost sharing produced about a 25 percent reduction in hospitalizations and in a related analysis, Lohr et al. (1986) concluded that the reduction in hospitalizations was not concentrated among those that were ineffective, but had the same mix of effective and ineffective procedures as in free care. Similarly, Siu et al (1986) concluded that the reduction was not concentrated in the inappropriate hospitalizations. That equally effective hospitalizations could be reduced by about a quarter through cost-sharing without any adverse health effects is a finding that should invite exceptional scrutiny.

The conventional interpretation of the RAND Experiment is that all the additional health care that fully insured consumers purchase represents “flat of the curve medicine.” Cost-sharing should, therefore, be imposed on insurance policies because it will save health care expenditures without any adverse health effects. My more skeptical interpretation is that some of this marginal health care—those equally effective hospitalizations that had been cut by about a quarter—cannot be forgone without incurring important adverse health consequences. That the RAND Experiment found differently is explained by participants who did not actually forgo hospitalizations, but instead dropped out of the Experiment when they were diagnosed with a condition that required a hospital procedure in order to avoid the cost-sharing associated with the hospitalization, which could have been as much as \$1,000 at the time (over \$4,000 in today’s dollars). If participants dropped out, they could revert back to their original insurance policies which, if employer-based, “tended to cover inpatient services well (frequently in full)...” (Newhouse et al., 1993, p. 402). This would explain the lack of diminished health effect for participants in the cost-sharing arm: participants in both the free care arm and cost-sharing arms were receiving the same level of care given their health conditions, thus their observed health status measures were similar. This reinterpretation, however, would have different implications for policy.

Direct information that would enable researchers to test whether this behavior had occurred, however, is lacking. The best information would be data on whether those who left the Experiment had a hospitalization after leaving, but data on utilization were only collected up to the time of the participant’s exit (Newhouse et al., 1993, p. 19). Alternatively, if data were available on the health status of those who dropped out *at the time of their leaving the Experiment*, we might expect some health response associated with the new diagnosis—for example, a diminished probability of an “excellent” self-reported health status—relative to those who did not drop from the cost-sharing plan. Unfortunately, health data were generally collected annually and it would have only been by chance that the health status would have been collected in the interval after a diagnosis and before receiving care. Therefore, we are left with trying to interpret what happened from the clues that the existing data and analysis provide.

Joseph Newhouse and his RAND Health Insurance Experiment colleagues present 5 objections to my interpretation of these clues in a paper published on Newhouse’s Harvard

University website as of October 25, 2007 (Newhouse et al., 2007). Below, I respond to these points in the order that they were presented.

Response to Newhouse et al. (2007)

1. Nyman falsely asserts that participants on the cost sharing plans facing a hospitalization had a financial incentive to drop out. This was likely true in the studies of plan switching he cites, but we designed the Experiment so that it would not be true. In fact, it was always in the financial interests of the Experiment's participants to remain enrolled because the Experiment made unconditional monthly side payments to the participants in the cost sharing plans that held them harmless against their worst case (the "Participation Incentive") and in addition paid them a lump sum if they completed the experiment (the "Completion Bonus.") Thus, there was no a priori reason to expect financially motivated attrition occurred. (We note in passing that all the Experimental cost sharing plans had a stop-loss feature that was at most \$1,000 for medical spending for the entire family for a year. As a result, even had there been no side payments, many individuals facing hospitalization would have been better off financially on the Experimental plan than on the plan to which they could return, many of which had no stop loss feature.)

As is argued in my article (Nyman, 2007), this participation incentive may have reduced the incentive to drop out from the cost-sharing plans, but did not eliminate it. First, the copayments for an anticipated hospitalization might have represented a cash flow problem that could be solved more easily simply by leaving the Experiment. Indeed, illness often exacerbates cash flow problems because of additional non-medical expenditures or a reduction in income. Second, Newhouse et al. (2007) implicitly assume that the value of the future payments from the participation incentive is constant for the individual. However, a person with a health issue is likely not to be as future-oriented as one without, therefore, she might discount the value of future payments for staying in the Experiment—perhaps severely, if the diagnosis is life threatening—compared with the more immediate issue of paying a large copayment. For these reasons, it is difficult to accept Newhouse et al's (2007) assertion that "there was no a priori reason to expect financially motivated attrition occurred."

Unfortunately, because of the uniqueness of the RAND Experiment, there is no comparable evidence (that I am aware of) to make the case that participation incentives of this type either work or do not work, and in science, an assertion in the absence of data is simply a hypothesis. The data we do have are from the RAND Experiment itself. The Experiment was designed so that the only differences between the free care and the cost sharing arms were financial differences, which in the cost-sharing arm included the copayments paid by the participant and the participation payments paid to the participant (v. neither in the free care arm). If a difference in behavior between the arms was observed, it would be reasonable to attribute that difference to differences in the financial treatment of participants since that was the sole experimental treatment (of this part of the Experiment). RAND found that attrition was significantly greater in the cost-sharing arms. Indeed, the percentage of participants who *voluntarily* exited from the cost-sharing was 16 times greater than the percentage that exited from the free care (Newhouse et al., 1993, p. 24). This is convincing evidence that this attrition

behavior was motivated by financing differences, or as I say in my paper, convincing evidence that the participation incentive simply did not work.

To my knowledge, no alternative explanation for this finding has been offered that does not raise selection issues. Selection was the reason that the RAND Experiment was implemented in the first place. Prior to the RAND Experiment, many observational studies had found that insurance was associated with increased consumption of health care, but there was always the question of whether these findings were due to ill people disproportionately purchasing (selecting) insurance. The RAND Experiment was intended to eliminate selection by randomly assigning participants to health plans that differed in financial coverage characteristics. That selection was the genesis for the Experiment magnifies the importance of any selection issues that may have been incorporated into the Experiment inadvertently. One such issue is the possibility that attrition at the point of diagnosis may have selected those with anticipated hospitalizations out of the cost sharing arm. Because of the substantial number of participants who dropped out and the compelling evidence that this attrition was financially motivated, this is a serious issue. To not have a benign explanation for this behavior is troubling, at least to this observer.

2. In his assertion that differential attrition can explain why cost sharing had negligible effects on health, Nyman ignores the fact that at the end of the Experiment we were able to collect health status data on 77 percent of those who had left the Experiment prematurely (85 percent of those who survived). These people are included in the reported general health results, so there is little reason to expect bias in the health status results.

If the cost-sharing arm includes data on the health status of drop outs after they had left the Experiment (for example, they left the Experiment on January 1, 1980, and responded to a health questionnaire on July 1, 1980) and in the meantime they had received a hospital procedure to treat their diagnosis just like they would have received if they had been in the free care arm (because the original policy to which they reverted had favorable coverage), then we would not expect any difference in health status measures between the free and cost-sharing arms. This was what the RAND Experiment found.

3. Similarly, our analysis gave us confidence that differential attrition did not materially bias the utilization results. Utilization while in the experiment by those who left was similar to those who did not leave (except for the one percent of participants on each plan who die, who experienced high end-of-life spending). We believe this finding greatly strengthens the likelihood that the utilization results are not an artifact of non-random attrition. Nyman, however, dismisses this argument with the assertion that individuals facing hospitalization would have withdrawn. Even ignoring the point made above that withdrawing would have left them worse off financially (because of side payments), his argument at best would only apply to anticipated hospitalizations; if one is suddenly hospitalized, for example with a heart attack, it would be too late to change insurance policies. And anticipated hospitalizations would likely have generated above average expenses in the period prior to hospitalization, as individuals sought care and physicians carried out diagnostic tests. We did not observe that.

In the absence of data on utilization after leaving the Experiment, Newhouse et al. (2007) suggest we could perhaps gain some insights on those who dropped out from the information that is available on utilization in the period before their leaving. Newhouse et al. (2007) hypothesize that *increased* utilization before leaving would represent evidence that those who dropped out were likely to have additional hospital utilization after leaving the Experiment. Thus, the lack of a relationship between attrition and utilization before leaving would indicate that the attrition was not motivated by anticipated hospitalizations.

While this is a reasonable hypothesis for some participants, equally reasonable would be the hypothesis that those who discovered a health issue during a routine physician visit may have dropped out in anticipation of cost sharing expenses from the diagnostic work-up and other outpatient services that might precede or be associated with a hospitalization. Attrition in the individual deductible arm was about as great as attrition in the other cost-sharing arms (and greater compared to the free care arm (Newhouse et al., 1993, p. 24)). The individual deductible covered inpatient services fully, but required a deductible for outpatient services, so the financial reason for leaving would be to avoid the payment of this deductible. (RAND researchers concluded that hospitalizations and outpatient services were complements [they were purchased jointly], based on the finding that the individual deductible reduced hospital utilization [Manning et al., 1987, p. 271]). The fact that more participants left this arm than left the free care arm implies that their financial position could be improved by dropping out.

Thus, drop outs may contain both those with increased consumption before leaving (from those in the individual deductible arm) and decreased consumption before leaving (from those in the copayment arm). Thus, from a theoretical perspective, it would be difficult to rely on a finding that there was no overall difference in utilization among stayers and leavers that attrition was benign.

4. Moreover, Nyman's speculation about a high degree of non-random attrition is contradicted in work that is unpublished but also posted on Newhouse's home page. In this work Manning, Duan, and Keeler carried out additional analyses on those who did not complete the Experiment. They concluded that there was in fact a modest amount of non-random attrition, but that its effects, if accounted for, would have left our conclusion that cost sharing reduced use unchanged.

Indeed, there is now empirical evidence available that differential utilization prior to leaving did occur that was systematically related to attrition. In the paper newly reported on Newhouse's web site, Manning, Duan and Keeler (1993) find that among all participants who dropped out (presumably these include the ones who left voluntarily, those who were terminated for health reasons such as being institutionalized in a nursing home, and those who were terminated for non-health reasons such as joining the military, from the Table in Newhouse et al., 1993, p. 24), inpatient use by the drop outs in the coinsurance arm (that is, the cost sharing arms, except for the individual deductible arm) was significantly lower by 35% (Manning, Duan, Keeler, 1993, p. 12). This is consistent with the interpretation that these participants anticipated the inpatient stay and dropped out to avoid the copayments, only to obtain the hospital care outside the Experiment. Those few who dropped out of free care had 88% more inpatient use, again consistent with strategic behavior and the anticipation of leaving a situation where full

coverage prevailed. Those in the individual deductible arm had insignificantly more hospitalizations in the period before leaving, but the insignificance could be explained by an offsetting mixture of (1) the additional hospitalizations by those who scheduled them early to take advantage of the full inpatient coverage and (2) the reduction of hospitalizations by those who delayed theirs so that they would avoid the copayments for the outpatient care that were associated with them. When the RAND data were re-analyzed to account for the differential utilization from attrition, the difference in hospitalizations between free care and all cost-sharing arms was about 19 percent, not 25 percent as has been reported (Manning, Duan, Keeler, 1993, p. 13). In other words, by not accounting for attrition, the difference in hospitalizations reported in the RAND findings was overstated by about 1/3.

Regarding the individual deductible group (for which inpatient care was free but outpatient care had a deductible), dropouts in all three arms (free care, coinsurance, and individual deductible) used less outpatient care before they left than those who stayed in the Experiment. The reduction was substantially greater for those in the coinsurance and individual deductible plans than in the free plan, but all were insignificant. While the sign is consistent with the interpretation that many in the cost-sharing arms were exiting the Experiment to avoid paying for the outpatient care that might be a precursor to a hospitalization, the insignificance might indicate a mixture of motivations among these dropouts.

The methodological approach used in this analysis predicted the inpatient (and outpatient) utilization by stayers with an econometric model, and then used that model to estimate utilization by drop outs, if they had behaved like stayers. Because illness is largely a random event, the econometric models on which these calculations were based could explain only 14% of the variation in hospitalizations (Manning, Duan, Keeler, 1993 p. 12) and 30% of the variation in copayment episodes (Manning, Duan, Keeler, 1993, p. 14). Thus, there is a large amount of variation in hospitalizations before leaving that is not explained. Moreover, the implications of these findings for the number of hospitalizations that occurred after leaving can only be suggested. Again, here we are attempting to make reasonable inferences regarding the behavior of those who left the Experiment after they left, based on data and analysis of how they behaved before leaving, because no data were collected on the utilization of drop outs after they had left the Experiment. Because of these uncertainties regarding the amount of utilization after leaving the Experiment but the clear evidence that strategic behavior was occurring, it is difficult to have confidence in the finding of little or no difference between the health of those in the free care arm and the cost-sharing arms, a finding based on the inclusion of health data from large portions of those who left. If the use of effective hospital care by those in the cost-sharing arms was sufficiently similar to the use of hospitalizations in the free arm, this would explain the lack of a health effect. Because of these uncertainties, it would hardly be appropriate to view as the “gold standard” the RAND conclusion that one could reduce equally effective hospitalizations by a quarter without adverse health effects.

To date, the orthodoxy of thinking surrounding the RAND Experiment has not encouraged much questioning of this or any of the results. I applaud Newhouse for making this additional work available for review because it reveals that attrition is truly an important threat to the confidence with which we can view this RAND conclusion.

5. Nyman's dismissal of the Experiment's result that cost sharing causally reduces utilization is further undermined by the enormous number of observational studies over many years, in many settings, with many different methodologies, that find utilization of medical services responds to relatively modest cost sharing (e.g., (Cherkin et al. 1989; Newhouse and Phelps 1976; Scitovsky and McCall 1977; Wedig 1988)). Recent work has focused on cost sharing for drugs and also finds an effect of differential copayment on use (e.g., (Goldman et al. 2004; Huskamp et al. 2005; Joyce et al. 2002; Landon et al. 2007)). Even the rate of emergency room visits respond to cost sharing, as it did in the Experimental results (e.g., (Hsu et al. 2006; O'Grady et al. 1985; Selby et al. 1996)). In these studies the income transfer effect of insurance than Nyman describes in the opening part of his comment does not have a material effect, because the amount of cost sharing is a small fraction of income.

I do not dispute that cost sharing reduces utilization. I think it does and agree that the non-experimental studies that Newhouse et al. (2007) refer to are generally convincing. The issue that I am addressing, however, is whether cost-sharing in the RAND Experiment actually produced a 25 percent reduction in equally effective hospitalizations and whether such a reduction would actually have no effect on health. Here, there is little corroborating evidence from the outside. Of course, much could have changed between RAND and recent practice, but a number of more recent studies (Tamblyn et al., 2001; Schoen et al., 2001; Piette et al., 2004; Heisler et al., 2004; Chandra et al., 2007) find that increased cost-sharing produces adverse health effects because of reduced utilization. Conversely but consistently, Miller, Vigdor and Manning (2004) conclude that the additional health care utilization from having insurance has positive health effects.

Other RAND findings also make the no-health-effects-from-reduced-hospitalizations conclusion suspect. First, the finding from the RAND Experiment that the poor and ill are more likely to be adversely affected by cost-sharing than higher income participants with illness (Newhouse et al., 1993, p. 208) may indicate that the poor were less likely to have favorable insurance coverage to revert to if they had dropped out (Waldmann, 2007). Therefore, they may not have had the same options as others and thus disproportionately stayed in the Experiment when diagnosed with an illness requiring hospitalization. Their decreased use is consistent with the conventional interpretation of the RAND results, but the adverse health effects are not. In this regard, it would be interesting to know whether the poor had the same level of insurance coverage upon entering the Experiment as did those from higher income groups.

Second, a puzzle of the Experiment was the absence of additional pain and worrisome symptoms among those in the cost sharing arm, compared to those in the free care arm. Indeed, participants in the free care arm were *more* likely to complain of pain and worrisome symptoms than those in the cost sharing arm (Newhouse et al., 1993, pp. 204-208). If participants were responding to cost-sharing by having medical symptoms or pain but not receiving care, one would expect (as the RAND researchers did) that there would be more unresolved symptoms and pain in that group. One interpretation could be more iatrogenic disease among those overusing medical care in the free care arm (Gruber, 2006), but an alternative explanation is that those in the cost-sharing arm dropped out when pain and worrisome symptoms occurred. If the follow-up health measures did not capture this completely because the leavers in the cost-sharing arms

were disproportionately ill and a portion of the responses from these drop-outs were missing (Newhouse et al., 1993, p. 24), this could explain the unexpected results.

Conclusion

The most important “lesson” of the RAND Experiment was that one could impose cost-sharing provisions on health insurance policies that reduced hospitalizations by about a quarter and have little or no effect on health. Reducing hospitalizations by 1/4 would save a considerable amount of health care expenditures, and because the RAND Experiment found that there are no important health consequences for the average participant as a result, this lesson has been a very attractive one and enormously influential on policymakers. Over the years, this finding has been used to justify a variety of proposals to use cost-sharing to reduce national health care expenditures.

The question that I raised in my paper in the *Journal of Health Politics, Policy and Law* (Nyman, 2007) is how confidently we can rely on this result. I have suggested that attrition might explain this finding, but there are other potential explanations. For example, some might argue that the health measures were not sufficiently sensitive to measure the health effect of these reduced hospitalizations. This is a reasonable explanation, but not one that I favor because of the broad spectrum of health status measures that were used in the RAND Experiment. Moreover, the reduction in health from being diagnosed with an illness but forgoing a hospitalization for a procedure that would effectively cure or improve that illness would likely have a major implications for health status that would register on at least one of the health status measures.

Attrition from the Experiment, however, seems to me to be the most likely candidate. First, a significant attrition differential was generated by the cost-sharing. Second, the number of *voluntary* drop outs alone is great enough to account for most of the difference in hospitalizations. Third, the new analysis presented by Newhouse on his web site (Manning, Duan, Keeler, 1993) shows that differences in utilization (before leaving) were systematically associated with attrition status. One compelling result is that those in the coinsurance arm who dropped out had 35 percent *fewer* hospitalizations before leaving than those in that arm who stayed in the Experiment. Fourth, based on a model that explains only 14 percent of the variation in hospitalizations, Manning, Duan and Keeler (1993) are still able to conclude that by not accounting for attrition, the effect of RAND cost-sharing on hospitalizations reported in the literature overstates the true effect by about 1/3. Fifth, a number of other findings in the Experiment—health effects from imposing cost sharing on the poor, the lack of an expected differential for painful or worrisome symptoms, and even the Lohr et al. (1986) result that cost-sharing did not disproportionately reduce the ineffective hospitalizations—are consistent with the attrition story. Sixth, contrary evidence from other studies on the relationship among cost-sharing, utilization and health status raises issues of whether this RAND Experiment conclusion can truly be considered the gold standard.

References

- Chandra A, Gruber J, McKnight R. Patient cost-sharing, hospitalization offsets, and the design of optimal health insurance for the elderly. NBER Working Paper 12972, Cambridge, MA: National Bureau of Economic Research, March 2007.
- Gruber J. *The Role of Consumer Copayments for Health Care: Lessons from the RAND Health Insurance Experiment and Beyond*. Menlo Park, CA: Kaiser Family Foundation, October 2006.
- Hiesler M, Langa KM, Eby EL, Fendrick AM, Kabeto MU, Piette JD. The health effects of restricting prescription medication use because of cost. *Medical Care* 2004;42(7):626-634.
- Lohr KN et al, Use of medical care in the Rand Health Insurance Experiment: Diagnosis- and service-specific analysis in a randomized controlled trial. *Medical Care* 1986;24(suppl.):S1-87.
- Manning WG et al. Health insurance and the demand for medical care: Evidence from a randomized experiment. *American Economic Review* 1987 77:251-77.
- Miller W, Vigdor ER, Manning WG. Covering the uninsured: What is it worth? *Health Affairs Web Exclusive*. 10.1377/hlthaff.w4.157, March 31, 2004.
- Newhouse et al. *Free For All? Lessons from the RAND Health Insurance Experiment*. Cambridge, MA: Harvard, 1993.
- Newhouse JP, Duan N, Keeler EB, Liebowitz A, Manning WG, Morris CN, Phelps CE, Rolph JE. Refusal and attrition in the RAND Health Insurance Experiment: A response to Nyman. Joseph Newhouse website, accessed October 25, 2007: http://www.hcp.med.harvard.edu/people/hcp_core_faculty/joseph_newhouse
- Nyman JA. American health policy: Cracks in the foundation. *Journal of Health Politics, Policy and Law* 2007;32(5):759-783.
- Piette JD, Wagner TH, Potter MB, Schillingere D. Health insurance status, cost-related medication underuse, and outcomes among diabetes patients in three systems of care. *Medical Care* 2004;42(2):102-109.
- Schoen MD, DiDomenico RJ, Connor SE, Dischler JE, Bauman JL. Impact of the cost of prescription drugs on clinical outcomes in indigent patients with heart disease. *Pharmacotherapy* 2001;21(12):1455-1463.
- Siu Al et al. Inappropriate use of hospitals in a randomized trial of health insurance plans. *New England Journal of Medicine* 1986;315:1259-66.
- Tamblyn R et al. Adverse events associated with prescription drug cost-sharing among poor and elderly persons. *JAMA* 2001;285(4):421-9.

Waldmann R. Personal communication. October 23, 2007.