

NBER WORKING PAPERS SERIES

ASSUMING THE CAN OPENER:
HEDONIC WAGE ESTIMATES AND THE VALUE OF LIFE

William T. Dickens

Working Paper No. 3446

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 1990

I would like to thank the Institute of Industrial Relations at Berkeley for generous research support. This paper is part of NBER's research program in Labor Studies. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

NBER Working Paper #3446
September 1990

ASSUMING THE CAN OPENER:
HEDONIC WAGE ESTIMATES AND THE VALUE OF LIFE

ABSTRACT

Although intuitively appealing, the use of hedonic wage estimates to determine people's willingness to pay to avoid the risk of fatal hazards is fraught with problems. The theoretical basis for such estimates are flawed in a number of important ways. The underlying behavioral model is wrong, there is imperfect information about job hazards, and labor markets do not look like the perfectly competitive model on which the theory depends for its conclusions. Further, there are many serious problems with the techniques used to estimate hedonic wage equations. This paper describes these problems.

Not surprisingly, these problems result in a wide range of results with respect to willingness to pay to avoid fatal hazards. It is argued that this wide range of results is not fully apparent in the literature because of the bias in publication towards positive as opposed to negative findings. The paper concludes that it is unlikely that economics has much to contribute to the public policy debate over the value of a life.

William T. Dickens
Department of Economics
University of California
Berkeley, California 94720

Assuming the Can Opener:
Hedonic Wage Estimates and the Value of Life
by William T. Dickens

It makes sense that dangerous jobs should pay more. Highly visible jobs such as test pilots or fire fighters do seem to pay more than comparable but less hazardous employment. Job evaluation systems which determine pay in many large firms always give "points" to jobs which expose workers to hazards, and pay increases with points. It is the common sense nature of this observation -- dangerous work should pay more -- which makes plausible economists' attempts to determine people's willingness to pay for less hazardous work using the hedonic wage technique. If people choose to be fire fighters or test pilots then certainly the wages they are offered must be adequate to induce them to take the risks. By comparing those wages to the wages of similar workers in less hazardous jobs we should get an estimate of how much job safety is worth to people. It seems only a small step to hedonic wage equation analysis in which the wages of workers are regressed on a long list of personal and job characteristics including the risk of a fatal accident. Why not interpret the coefficient on job risk as what workers are willing to pay to avoid a unit of risk?

Why not? Because there are three fundamental problems with doing so. First, the behavioral foundation of the economic theory of hedonic estimation is inappropriate. People don't

behave toward risk the way economists assume that they do. Second, even if people behaved the way economists assumed, the additional assumptions necessary for hedonic estimates to yield willingness to pay are so far fetched that the estimates cannot be taken seriously -- people do not have perfect information about job safety and labor markets are not perfectly competitive.

Finally, even if there were no problem with the theory, there are many serious problems with the data and with the techniques used to estimate hedonic wage equations. Any one of these three problems should make one wary of the results of these studies. Taken together they should be seen as seriously challenging the usefulness of the enterprise.

If there are so many problems with hedonic studies, why do they yield reasonable results? My answer is that they don't -- they only appear to. A wide range of results are possible. However, the process of publication places a screen between the reader and the researcher which results in the appearance of some uniformity in the results. The uniformity reflects the prior opinions of authors and editors far more than the willingness to pay of labor market participants.

Hedonic wage estimation is only one of several methods typically used to deduce people's willingness to pay to avoid the risk of a fatal accident. Many of the criticisms raised below apply with equal force to the consumer and contingent valuation studies. The criticisms of the behavioral model are an example. The lack of perfect information is also a problem in the product

market and affects the usefulness of consumer valuation studies. One could conclude that since these methods have fewer flaws than hedonic wage studies they represent a superior method. On the other hand, one could also conclude that the remaining problems are still serious and render useless the results of these analyses as well.

What Hedonic Estimation Assumes

The jump from "hazardous jobs pay more" to the conclusion that the coefficient on fatal accident risk in a hedonic wage equation measures workers' willingness to pay to reduce risk is a big one. While what immediately comes to mind when we think about risky jobs are a handful of very visible jobs -- such as fire fighters, pilots, police, or high-rise construction workers -- these represent only a small fraction of the jobs in the economy. It is not the contrast between these jobs and other jobs that drives hedonic wage equation estimates. Rather it is the relatively marginal differences in safety between people working in manufacturing, trade, and services and the differences between industries within these broad groups.

For this to work people have to be able to perceive and make rational judgements about very small differences in the probability of a fatal accident. This is expecting a lot of people's cognitive abilities and their knowledge of work hazards -- perfect information, perception and cognition.

Beyond these problems, the economist interpreting hedonic wage equations as reflecting people's willingness to pay is also assuming that the labor market is highly competitive.

Specifically, he or she assumes that employers only pay the minimum necessary to attract workers to a job. In so doing the economist ignores unions, unemployment, and other labor market imperfections.

Finally, the statistical techniques used to estimate hedonic wage equations require the assumption that all determinants of an individual's wage are either observed by the analyst or are uncorrelated with the measure of risk. Both economic theory and problems with the risk data guarantee that this assumption will not hold. The next several sections explore these problems in more detail.

The Behavioral Model is Wrong

Economists typically assume that people maximize their own welfare. When outcomes are uncertain economists assume that people take this into account and act to maximize their expected welfare. This gives a very neat and precisely defined model of how people should behave. Sufficiently well defined that a number of psychologists and economists have been able to test it in a wide variety of circumstances inside and outside the laboratory. The unescapable conclusion is that people don't behave the way economists assume that they do.

In particular, these studies conclude that people are at their worst making individual decisions involving low probability events when they have little or no chance to learn from their mistakes. Job decisions involve precisely these elements if they require workers to consider the probability of a fatal accident. Many different types of anomalous behavior have been identified. Dickens (1985, 1986) discusses the wide range of findings in the psychological decision theory literature relevant to job and product safety. A complete review would require a much longer paper. This paper reviews three of the most important problems.

When making decisions among choices with many aspects, people tend to simplify the decision by not paying attention to certain aspects of the choices and limiting their attention to a subset of the characteristics of each choice. In particular, people tend not to pay attention to aspects of a choice which have only a very low probability of occurring. This tendency has been used to explain why people don't buy subsidized flood and earthquake insurance (Kunruther, 1978; Kahneman and Tversky, 1979). It should also apply to job choice. For the overwhelming majority of jobs, the risk of a fatal accident is very low. From the evidence on how people make other types of decisions, it seems likely that most people ignore the probability of a fatal on-the-job injury in choosing jobs.

If many people do not pay attention to such risks, then employers have only limited incentives to pay compensating differences and measured differences will not reflect willingness

to pay to avoid those risks. If compensating differences appear at all, they may reflect employers' costs of providing safety rather than employees' preferences (Dickens, 1985).

When people do pay attention to low probability events in making decisions, they tend to behave as if they thought such occurrences were much more likely than they are but to discriminate less between low probability outcomes than they should¹. This could lead workers to demand a much higher compensating difference when comparing a hazardous job with one which was safe, or too little when compared to another job with less but still positive risk.

These two problems could be the result of limited rationality -- people could be trying to behave in the manner economists assume they do, but they may not have the cognitive resources to do so. A third problem suggests that people just don't think in expected welfare terms. Many studies have found that the way a problem is framed affects the value people place on an outcome. For example, people will pay less to improve their circumstances than they are willing to pay to prevent what they view as a loss (Kahneman and Tversky 1986). Thus the question of willingness to pay to avoid risk may not be one that has a unique answer in most people's minds and the concept of a single value of life for a particular individual -- even holding his or her income constant -- may be inappropriate.

Viscusi (1984) has looked at the rationality of some aspects of worker behavior toward risks. In particular he shows that

workers' responses to being told that they will be working with certain hazardous substances can be modeled as rational Bayesian decision making -- workers respond to the information by adjusting their assessment of the job risk and the amount of compensation they say they would require to stay on the job. However, it is hard to imagine a theory of behavior for which this would not be true. Further, nothing in this work, or any other work, directly addresses the specific concerns raised by the literatures discussed above. Nothing suggests that editing isn't important in job choice or that workers don't over-weight low probability events. Nothing addresses the issue of framing. Dickens (1985) describes how all existing evidence on labor markets and job safety can be easily reconciled with the psychological model discussed above.

The Assumptions of the Economic Model Don't Hold

As noted above, the economic model behind the estimation of willingness to pay to avoid risk assumes that workers not only are rational about such decisions, but that they are perfectly informed and that the labor market is perfectly competitive. The effects of deviations from these assumptions have not been thoroughly explored, but to the extent that they have it appears that they are crucial to the interpretation of hedonic studies.

Take first the problem of perfect information. No one would assert that there aren't substantial information problems that affect peoples' job choices. Viscusi's work acknowledges and

documents these problems (Viscusi 1979, 1984). He finds that quit rates are higher in more dangerous jobs. This suggests that many workers don't know how dangerous a job is when they take it but that they learn over time. Their quitting also suggests that whatever wage differences they receive are inadequate to compensate for their jobs' true risks. I know of no model that specifically considers the size of compensating differentials in the presence of imperfect information. My own work on this topic suggests that with rational workers, differences will reflect a combination of worker's preferences and firms' costs of providing safety (Dickens 1984b). Thus this problem alone would probably result in hedonic estimates understating peoples' true willingness to pay.

No such statement can be made about the direction of the bias if labor markets are not competitive. Although they are a shrinking fraction of the labor force, unions still bargain wages for about 15% of all workers. A much larger number of unorganized workers receive wages similar to union wages from employers who wish to avoid unionization. Foulkes (1980) reports interviews with compensation specialists at major non-union firms. He concludes that pay received by union workers at comparable firms is the major determinant of pay at these firms. A number of recent studies suggest that the payment of super-competitive wages is the norm rather than the exception². The payment of high wages by a large number of firms is consistent with significant persistent unemployment.

In a competitive market a firm which offers hazardous work must compensate workers for the full amount of their perceived cost of exposure to the hazard. However, a firm which is already paying above the market wage need not pay that additional cost to attract workers. On the other hand, a union or a worker may be able to use the presence of a hazard to bargain a larger wage increase if, for example, the cost to the employer of eliminating the hazard is much larger than the value to the employee(s)³. Thus even if workers are rational, one cannot say a priori whether hedonic estimates of willingness to pay to avoid job risk are under- or over-estimates of true willingness to pay.

The Assumptions of the Econometric Methods Aren't Satisfied

Even if the economic and behavioral theory of hedonic estimation was sound, existing estimates of the value of life based on labor market data fall far short of providing believable estimates for another reason. There are at least three major problems with the statistical analysis done to arrive at these numbers.

When estimating the determinants of wages using multiple regression analysis the effects of any one variable are identifiable only to the extent that any determinants which are not observed by the analyst are uncorrelated with the variable of interest⁴. However, the economic theory of hedonic wage equations insures there are omitted variables that are correlated with job safety (Brown, 1986; Kahn and Lang, 1988). The omitted

variable problem is much more severe here than in other types of economic problems.

In general we know that it is impossible to control for all aspects of individual earnings potential in a wage regression. At best, human capital equations explain about half of the variance of individual earnings. To the extent that earnings potential is not captured by included variables and that job safety is a normal good -- those with more income want safer jobs -- unobserved earnings potential will be negatively correlated with job risks and will bias the estimates of willingness to pay to avoid risk downwards.

A further problem is other undesirable working conditions closely associated with fatal job risk but not included in the analysis. Dirty jobs, fast paced jobs, jobs involving heavy lifting, exposure to extreme temperatures or noxious chemicals are also more likely to be dangerous jobs. This is just a partial list of such characteristics. Most studies cannot control for any of these characteristics. None can control for all of them. Workers may not have a very good idea of what the risk of a fatal accident is, but they may have a very good idea of how unpleasant the job conditions are. In this case the coefficient on risk of a fatal accident would be measuring worker's aversion to these other characteristics as well as their willingness to pay to avoid risk. The importance of risk would tend to be overestimated. In the extreme, estimated compensation

for hazardous work could reflect only the compensation for other undesirable attributes of hazardous jobs.

A second major econometric problem is the robustness of the estimates of willingness to pay. It is not uncommon in economics to find that using different data sets or different sets of control variables in a multiple regression leads to a variety of answers to the same question. Although most published results find positive and significant compensating differentials there are exceptions. Viscusi (1980), Dorsey (1983), and Dickens (1984a) all find negative returns to fatal job risk in the non-union sector. Ruser (198?) finds no evidence of risk premiums and Leigh (1987, forthcoming) finds no evidence of compensating differences when using occupational risk data.

Are these findings of small or negative differentials an aberration or do they reflect a fundamental problem with the estimates? Can they be ignored since the bulk of published studies show substantial returns to fatal job risk?

No data set tells us precisely what an individual's risk of a fatal job accident is. As a result, what most studies have done is to match data on industry or occupational fatalities with the occupation or industry of individual workers and the wages they receive. There are several problems with this approach. One of the biggest is that industries have many characteristics besides their fatal job risk. In a study of the relation between wages and industry characteristics, Dickens and Katz (1987) ran 432 different regression specifications -- varying the sample and

the control variables used -- to study the robustness of those relations. All specifications examined were equivalent to or very similar to ones found in the literature on wages and industry characteristics. We found that of 22 industry characteristics only 2 had the same sign in all specifications. Although fatal hazard risk was not considered, risk of injury was. Injury risk was one of the variables with the least consistent relation to wages. In general these results suggest that the effects of different industry characteristics on wages are difficult, if not impossible, to sort out.

The matching of average industry hazard rates with individual data to estimate hedonic wage equations causes two other problem as well. Even in very large industries the number of fatal accidents per year is relatively small. Consequently, the probability of a fatal accident is inaccurately measured in many existing data sources. This leads to a standard errors-in-variables problem which tends to cause an underestimation of the effect. Techniques exist for correcting this problem (Dickens and Ross 1984), but to the best of my knowledge no hedonic wage study has ever undertaken to do so.⁵

The final problem is caused by matching industry level data with individual data. Standard computer packages compute significance levels and standard errors on the assumption that there are as many unique observations as there are individuals. However, with the independent variable observed only at the industry level there may be, in effect, far fewer observations.

As a result the precision of hedonic wage estimates are probably greatly overstated (Moulton, 1986). If standard errors were computed correctly the positive effects which have been found might not be statistically significant.

Estimates of Hedonic Damages and the Sociology of Economics

If the relationship between job risk and wages is as fragile as the above discussion suggests, why do so many published studies find large significant results in a range that seems to most to be quite reasonable? From my own experience estimating these type of relations I don't believe that the published results are representative of the data. Instead I think they are more reflective of the prejudices of the studies' authors, referees and journal editors.

"Dog bites man" is not news. The social science equivalent of "dog bites man" is "existing techniques and data are inadequate to answer the question." Scientific journals want to publish interesting articles so editors avoid publishing negative findings. Researchers know this and will frequently abandon research programs that are failing to produce interesting results.

If a wide range of results are possible, depending on the specification one chooses, then with no artifice or guile it is possible for the journal editing and refereeing process to result in the publication of only those results consistent with what authors, referees, and editors feel are reasonable results.

Suppose several different researchers try several different specifications. Those who get unbelievable results attribute their findings to the sorts of problems discussed above. A few may try to publish their findings, but editors and referees agree that the article has major flaws (all do) that account for the poor results. The results are judged uninteresting and are not published. On the other hand, those studies which find believable results have an easier time getting through the editing and refereeing process -- if a referee or editor raises a question about the believability of the method they then face the question, "if the method is flawed why are the results believable?" Thus the believable results have a chance of getting into print and some do.

Not surprisingly, those negative studies which have been published have appeared in conference proceedings and less prestigious journals than the studies showing positive "believable" findings.

Conclusion

Although intuitively appealing, the use of hedonic wage estimates to determine people's willingness to pay to avoid the risk of fatal hazards is fraught with problems. The theoretical basis for such estimates are flawed in a number of important ways. The underlying behavioral model is wrong, there is imperfect information about job hazards, and labor markets do not look like the perfectly competitive model on which the theory

depends for its conclusions. Further, there are many serious problems with the techniques used to estimate hedonic wage equations.

Not surprisingly, these problems result in a wide range of results with respect to willingness to pay to avoid fatal hazards. This wide range of results is not fully apparent in the literature because of the bias in publication towards positive as opposed to negative findings.

What is the cumulative effect of these problems? Can one say whether the standard estimates are too low or too high? As noted above, it is not clear that the question even makes sense. The psychological literature suggests that people don't think in terms of a unique value for avoiding a certain level of a hazard. Rather their willingness to pay is highly dependent on the specific circumstances -- how the problem is framed.

Ignoring this problem, it is still impossible to say whether estimates are biased up or down. Information problems and most of the problems with less than perfect rationality would lead to underestimates of willingness to pay, as would the errors-in-variables problem. However, the lack of competition in labor markets, and particularly the problems of omitted job characteristics, may lead to over-estimates. I suspect that true compensating differentials would be very hard to find in the data for all of the reasons cited above, but that other industry attributes associated with fatal job risk lead to a spurious correlation between fatal risk and wages.

My own work (Dickens, 1984a) suggests that it is mainly the difference in wages between manufacturing industries and other industries which drives the correlation between wages and fatal job risk in the non-union sector. There are many other explanations for why manufacturing pays more than other industries besides greater risk of fatal injuries.

Many of the criticisms raised above -- particularly those of the behavioral model -- apply to other methods for determining how much individuals value safety. It is not clear that any of these methods represent a superior approach to hedonic wage estimates. It may be that economics simply has nothing useful to say about this issue in which case policy makers and juries should be left to their own resources of common sense and introspection to answer the question of how much a life is worth.

References

Brown, Charles, "Equalizing Differences in the Labor Market," Quarterly Journal of Economics, (February 1980) p113-134.

Dickens, W.T., "Differences Between Risk Premiums in Union and Nonunion Wages and the Case for Occupational Safety Regulation," American Economic Review, Papers and Proceedings, Vol. 74 #2 (May 1984a) p320-323.

----- "Occupational Safety and Health Regulation and Economic Theory," in William Darity Jr. ed. Labor Economics: Modern Views, Hingham MA: Kluwer-Nijhoff, (1984b), p133-173.

----- "Occupational Safety and Health Regulation and Irrational Behavior: A Preliminary Analysis," in John D. Worrall and David Appel, eds., Workers' Compensation Benefits: Adequacy, Equity and Efficiency, ILR Press, (1985) p19-40.

----- "Safety Regulation and 'Irrational' Behavior," in Benjamin Gilad and Stanley Kaish, eds., Handbook of Behavioral Economics, JAI Press, 1986.

Dickens, W.T. and Lawrence Katz, "Inter-Industry Wage Differences and Industry Characteristics," in Kevin Lang and Jonathan Leonard, eds. Unemployment and the Structure of Labor Markets, Basil-Blackwell, 1987.

----- "Inter-Industry Wage Differences and Theories of Wage Determination," NBER working paper #2271 (1988).

Dickens, W.T. and Brian Ross, "Consistent Estimation Using Data from More than One Sample," NBER Technical paper #33 (1984).

Dorsey, Stuart, "Employment Hazards and Fringe Benefits: Further Tests for Compensating Differentials," In John D. Worall ed., Safety and the Workforce: Incentives and Disincentives in Workers' Compensation. Ithaca: New York State School of Industrial Relations, 1983.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decisions Under Risk," Econometrica, Vol. 47, #2 (March 1979), p263-291.

-----, "Rational Choice and the Framing of Decisions," Journal of Business, Vol. 59, #4 (Oct. 1986), p s251-78.

Kruger, Alan and Lawrence Summers, "Reflections on the Inter-Industry Wage Structure," in Kevin Lang and Jonathan Leonard, eds. Unemployment and the Structure of Labor Markets, Basil-Blackwell, 1987.

-----, "Inter-Industry Wage Differences and Efficiency Wages," Econometrica, (1988)

Kunreuther, Howard, Disaster Insurance Protection: Public Policy Lessons, New York: Wiley, 1978.

Lattimore, P.K., J. R. Baker and A. Dryden Witte, "The Influence of Probability on Risky Choice: A Parametric Examination," NBER Technical paper #81.

Leigh, J. Paul, "Gender, Firm Size, Industry and Estimates of the Value of Life," Journal of Health Economics, Vol. 6 (1987):255-273.

-----, "Data on Occupational Fatalities Fail to Provide Evidence for Compensating Wages," Industrial Relations (forthcoming).

Viscusi, W.K., Employment Hazards: An Investigation of Market Performance, Cambridge: Harvard University Press, 1979.

-----, "Union, Labor Market Structure, and the Welfare Implications of the Quality of Work," Journal of Labor Research, Vol. 1, (Spring 1980), p175-192.

Viscusi, W. K. and Charles O'Connor, "Adaptive Responses to Chemical Labeling: Are Workers Bayesian Decision Makers?" American Economic Review, Vol 74 (December 1984), p942-56.

Notes

1. Khaneman and Tversky (1979) discuss the evidence for this phenomena. A recent study of the exact nature of the relation is Lattimore, Baker and Witte (1989).
2. See Dickens and Katz (1987,1988), Krueger and Summers (1987,1988), and Katz and Summers (1989) for reviews of this evidence.
3. In my 1984b paper on hedonic wage equations I demonstrate that increasing job hazards can be associated with either larger or smaller premiums when wages are determined by a Nash bargain as compared to a competitive market.
4. Actually the conditions are more restrictive than this but problems caused by correlations between left-out variables and other included variables are second order and not likely to be important.
5. Some recent studies have used five year averages. This substantially reduces the measurement error.