The Quality of Jobs

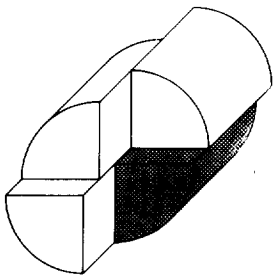

Employer-provided benefits: employer cost versus employee value

Cash-equivalent value is one approach to measuring employees' value of noncash benefits; more data and research are needed, however, to resolve complex methodological issues regarding this approach

Melissa Famulari
and
Marilyn E. Manser

Employers compensate their employees not only with cash, but also with noncash payments. The latter, sometimes called "in-kind" or "fringe" benefits, include some that are legally required, such as Social Security, workers' compensation, and unemployment insurance, and some that are not, such as paid leave, health and life insurance, and pensions. This article discusses the effort of economists to measure the value individuals place on noncash payments.[1]

Economists have developed the concept of "cash-equivalent value" to measure the value of noncash benefits to an individual.[2] A person's cash-equivalent value is the least amount of money he or she would be willing to accept in exchange for not receiving particular noncash goods. When applied to an employer-provided benefit, the cash-equivalent value is the minimum amount of additional cash compensation the worker will accept in lieu of receiving the benefit. Although some estimates of cash-equivalent value for Government-provided in-kind benefits such as food stamps, public and subsidized housing, medicaid, and medicare exist, little has been done to quantify employer-provided benefits, primarily because of a lack of data. Moreover, even if data were available, a variety of problems have made it difficult to implement the cash-equivalent value approach.[3]

Information on the employer's cost of providing the benefit is readily available through the Bureau of Labor Statistics' Employment Cost Index (ECI) program. To what extent does employer cost approximate employee value? We believe that there are various biases associated with using employer cost as the measure of the employee's value of employer-provided benefits. These biases can result in misleading conclusions.

This article (1) outlines the cash-equivalent value approach to measuring value and points out the relationship of employer cost to cash-equivalent value, (2) describes the three most commonly utilized techniques to estimate cash-equivalent value and examines the data requirements and limitations of each technique, and (3) discusses both the types of studies where the biases do not appear to distort conclusions and studies where we feel the use of employer cost as a measure of employee value would result in highly misleading conclusions. While some of this material is quite technical, we have endeavored to keep the discussion as general as possible.[4]

Employee value

Perhaps the simplest approach to estimating value is that which holds value to be the employer's cost of a given benefit. A more sophis-

Melissa Famulari is an economist in the Office of Economic Research, Bureau of Labor Statistics. Marilyn E. Manser is assistant commissioner of the same office.

ticated view is the *funds-released* approach, which maintains that value can be measured as the amount of money an individual would have spent to acquire a certain good in the absence of being provided the benefit associated with that good. In contrast to this is the *market-value* approach, in which value is the amount the individual would have paid for the benefit if he or she had purchased the *specific* amount the employer provided (as opposed to the amount the individual would have chosen at the existing market price). Finally, there is the *cash-equivalent* approach, wherein value is the least amount of money an individual would be willing to accept in exchange for not receiving a certain benefit. This last approach is the way most economists define value.[5]

Description of the theory. If an individual is given some noncash benefit by his or her employer, such as free or subsidized meals, pensions, or health insurance, and nothing else changes, plainly the individual will be better off than before the receipt of the noncash good. But *how much* better off? The cash-equivalent value approach to determining the change in individual well-being is fairly straightforward: we ask, "What is the minimum amount of additional cash compensation an individual would require to become just as well off as that individual would be if he or she received the noncash good?"[6] The cash-equivalent value is then the amount of cash compensation which makes the individual indifferent between getting the benefit and no cash and getting the cash and no benefit.

Sources of difference between employer cost and employee value. In many ways, it would seem that the relationship between employer cost and employee value is fairly straightforward. Any employer-provided benefits other than legally required ones result from employer-employee contracting, directly for unionized, and less directly for nonunionized, workers. Except for legally mandated benefits, employers can compensate employees with either wages or noncash benefits, and because both cost the employer the same amount, in general the employer has no incentive to prefer one form of compensation over the other.[7] Thus, in a perfectly competitive situation, in the absence of government intervention, taxes, and other institutional restrictions, the marginal worker (the last worker convinced by the total compensation package to accept the job) would be expected to value the last dollar of each benefit type equal to another dollar of money wages, or he or she could be made better off at no cost to the employer. Consequently, in such a situation, the cost of the benefit provided is a measure of its value.

Despite this apparent consonance, there are a number of reasons why employer cost can diverge from employee value. Employer cost for a noncash benefit that is not legally required can diverge from employee value because (1) the benefit is not subject to personal income taxes, (2) the benefit is provided uniformly to large groups of employees in a firm, and (3) the employer's cost of providing the benefit may be lower than the market price of the benefit to the employee. Let us consider each of these in turn. (Several graphical representations of the disparity that may arise between employer cost and employee value are presented in the appendix.)

1. When the benefit is not subject to income taxes, the equality of employer cost and employee value breaks down even for the marginal worker. Specifically, in a world with taxes, the marginal worker would be expected to consume noncash benefits up to the point where the marginal value of another dollar of benefits equals the after-tax value of another dollar of money wages. That is, the marginal worker, whose marginal tax rate is, say, t, would need to receive $1/(1-t)$ dollars of pre-tax income in order to get a dollar of after-tax income. (The amount $t/(1-t)$ is paid in taxes, leaving $[1/(1-t)] - [t/(1-t)] =$ one dollar.) More of the benefit will be consumed than if this differential tax treatment did not exist, and as a result, employer cost will overstate the cash-equivalent value of the benefit. (One common explanation for the growth of employee benefits has been the increasing tax rates faced by typical workers.) An implication of the differential tax treatment is that higher income workers will place a higher value on noncash benefits than will lower income workers, because the tax rate of the former exceeds that of the latter.

Various biases associated with using employer cost as the measure of employee value can result in misleading conclusions.

2. Historically, firms have typically provided benefits to their work forces as a whole, instead of tailoring them to the preferences of individual employees.[8] In addition, laws mandating that personal income tax advantages be available only for benefits which do not favor higher paid workers provide incentives for a more uniform provision of benefits to all employees in a firm. If there are costs associated with changing jobs and if employees are not perfect substitutes for one another in production, then uniform provision of benefits drives a wedge between employer cost and employee value for at least some individuals.[9] Under these conditions, the employer would be expected to provide benefits in accordance with the preferences of the "median" worker.[10]

Even in the absence of differential tax rates, higher income workers would be expected to demand more of any "normal" good (a good which people want more of as their income increases) than would lower income workers. In support of this proposition, Steven A. Woodbury estimated an income elasticity for noncash benefits which is greater than one; that is, a 1-percent increase in income leads to a greater-than-1-percent increase in the demand for noncash benefits.[11] Consider the class of benefits whose provision does not typically vary with employee income (for example, health insurance, child care, Christmas bonuses, and parking). Because employer cost will be the same regardless of employee income, the benefit ratio, that is, the ratio of employee value to employer cost, would be expected to be higher for higher income workers than for lower income workers in a firm. By contrast, for benefits which are provided in amounts proportional to income, such as life insurance and pensions, less of a difference in benefit ratios among workers would be expected.

When a benefit is not subject to income taxes, the equality of employer cost and employee value breaks down.

In general, benefit ratios are also expected to vary by some demographic factors, particularly family status. For instance, two-earner families that receive largely duplicative health insurance policies would place a relatively low value on one of them, single individuals may place a low value on life insurance policies, and so on. Thus, assuming that employee benefits are valued equally by all households—even at a given income level—may severely distort comparisons of well-being among households.[12]

As of 1978, cafeteria plans—that is, plans whereby workers choose among a "menu" of benefit options—could qualify for tax-exempt status. Depending upon the particular choices the employees have, these plans allow the provision of benefits to vary among the workers of a firm. That individuals do choose differently when given a choice is additional evidence that uniformity is a factor in driving a wedge between employer cost and employee value. We may assume that, as more options are given, the amount of the benefit the employer provides approaches the amount the employee would choose. Anecdotal evidence that employees do choose differently when given a choice is presented in the article "Flexible benefit plans: employees who have a choice," by Joseph Meisenheimer II and William Wiatrowski, elsewhere in this issue.

3. Ignoring taxes and issues relating to the uniform provision of benefits, an employee's marginal value of a noncash benefit will be set equal to the employer's marginal cost of providing the benefit, everything else being equal. If the employer's marginal cost is lower than the employee's market price, then the benefit will be "overprovided" relative to the amount the employee would choose at the market price. Thus, employer cost will provide a lower bound on employee value, and the amount provided by the employer times the price the employee would pay in the market will provide an upper bound.

The cost to the employer may be less than the purchase price to the employee for three reasons.[13] First, employers are often able to take advantage of discounts sellers offer for bulk purchases. Second, sellers are willing to provide benefits to groups of people at a cheaper rate than to individuals when there are adverse selection problems. Adverse selection occurs when, for example, the workers in the poorest health are the ones who want to purchase the most health insurance. By selling benefits to a group of workers, sellers can mitigate adverse selection. Finally, employers may prefer providing more of a given type of benefit than is demanded in order to reduce turnover, maintain a healthier and more productive work force, or attain another, similar objective. Here, a more inclusive measure of employer cost—one which "nets out" the gain accruing to the firm in providing the benefit—would result in an employer cost that is less than the market price.

The case of legally required employer-provided benefits is different from that of nonrequired benefits. The difference between the two is that employers and employees can negotiate about which nonrequired benefits will be provided in what amounts, whereas quantities of legally required benefits may be arbitrarily set with regard to employee values. Even given that employees as a group vote for legislators who will enact desired changes in the provision of mandated benefits, some voters who pay taxes are neither workers nor participants in the labor market, so median voter results may not apply.[14] In addition, the aforementioned problem concerning the uniform provision of benefits is exacerbated in the case of legally required benefits. Legally required benefits are provided uniformly to the work force as a whole and not just to groups of workers within a firm. Further, because unemployment insurance and workers' compensation are not fully experience rated, workers in some industries or firms will place a relatively higher value on them than will other workers.[15] As a result, while there may be some cases in which it is reasonable to use employer cost as a measure of the typical employee's value for benefits that are not legally required, it seems quite unrealistic to assume that employer cost of legally mandated benefits is a reasonable approximation of employee value.

Estimating cash-equivalent value

Given the many discrepancies between the employer's cost and the employee's value of a certain benefit, just how difficult is it to get a measure of the cash-equivalent value of the benefit?[16] The discussion that follows shows that estimation of cash-equivalent values is in general quite difficult, both because the data requirements are so extensive and because of the complex issues involved in the actual estimation, even if the data were available. Certainly, if more high-quality data existed, more research would be done on the methodological issues. Three approaches are currently the most frequently utilized:[17]

1. *Utility-Based Estimates.* Researchers have estimated recipient values by assuming some functional form for utility.[18] This sets a particular functional form for the demands for goods. Theoretically, because we can observe quantity demanded and price, and economic theory suggests which variables affect the demand for goods (although demand is also affected by factors influencing utility functions across individuals—something about which economic theory has nothing to say), demands for goods can be estimated. After estimating a particular demand system, researchers use the parameter estimates to compare the costs of achieving levels of utility with and without a given noncash benefit. This permits the calculation of the cash-equivalent value of the benefit. Ideally, the data needed to support such a technique include information on prices, wages, amounts of leisure and goods consumed, and characteristics of the benefit package. Although to our knowledge, there are no data sets with all the desired data, there are some studies that employ the utility-based technique to examine recipient values of Government transfer programs.[19]

There are a number of methodological problems with utility-based estimates of cash-equivalent value. Probably the foremost is that it is computationally difficult to estimate a demand system (much less, get all the data) for all the goods people demand. Somewhat questionable assumptions must be made in order for the estimates to be valid, for example, that (1) today's consumption is unrelated to both past and future consumption, (2) a reasonable functional form for utility has been chosen, and (3) utility functions among individuals are the same, at least within demographic subgroups.

2. *Survey Approach.* This technique involves asking employees directly about their willingness to pay for various noncash benefits. Questions like "What is the maximum amount you would be willing to pay to receive this benefit?" are posed to individuals who do not have the benefit in question. Support for this approach in regard to valuing public goods such as environmental quality exists in the literature,[20] because study results are both replicable and logically consistent with the predictions of demand theory. In addition, evidence that the survey approach yields the predicted magnitudes relative to hedonic wage equation results (see next) is also in the literature.[21]

The two chief problems with the survey approach are that (1) estimates of value are based upon hypothetical as opposed to actual choices and (2) empirically, there appears to be a significant downward bias in people's stated values—estimates of 50–67 percent exist in the literature.[22]

3. *The Hedonic Approach.* The theory behind the hedonic approach, which was popularized by Sherwin Rosen in 1974, is that variation in the observed mix of benefits and cash compensation offered by employers competing for workers having the same productivity is the result of the different tastes for benefits of those workers and the differential ability of employers to provide those benefits. In theory, ignoring institutional features discussed earlier, the amount of wages given up to obtain a specified amount of a noncash benefit is a measure of both the marginal value of the benefit to a worker who accepted the wage-benefit compensation package and the marginal cost to the firm in providing the benefit.[23] The simplest application of the hedonic approach would estimate a regression equation relating the wage to the amount of a particular benefit offered, all else being equal.[24] However, movements along the function given thereby reflect both differences in worker tastes for benefits and the firm's ability to provide the benefits. As a result, such movements do not, in general, provide a measure of the change in employee value for significant changes in the amount of the benefit provided.[25]

Under certain conditions, uniform provision of benefits drives a wedge between employer cost and employee value.

The hedonic wage equation tells one very little about either the demand for or supply of benefits; rather, it provides an estimate of a single point on compensated demand and supply functions. However, identifying the underlying compensated demands could provide an estimate of the cash-equivalent value of the benefits. Rosen has suggested a second stage to the hedonic method which would enable a researcher to identify the said cash-equivalent values.[26]

There are difficulties, however, even in the first-stage estimation—that is, estimating the wage differential associated with the differential provision of benefits.[27] First, because the level of benefits provided is often based upon the amount of wages paid to the individual, statistical complexities arise.[28] Second, researchers must assume that individuals were able to obtain the desired amount of the particular benefit being studied, rather than having had to choose among a limited number of packages of wages and benefits.[29] Finally, because no particular shape of the hedonic wage function is specified by theory, researchers need to allow for variety in their empirical estimates.

In the second-stage estimation, identification of the underlying compensated supply and demand parameters has proven to be far more complex than originally anticipated by Rosen.[30] Data requirements are significant for the proper implementation of the proposed technique, and to our knowledge, no empirical estimates of cash-equivalent values for employee benefits exist in the economics literature. By contrast, estimates of compensated demands for clean air, housing amenities, neighborhood characteristics, and noise from hedonic price equations do exist in the literature.[31]

Conclusions

In examining the issues that arise in obtaining employee values of employer-provided benefits, we conclude that employer cost is limited as a measure of employee value. For some purposes, however, using employer cost to proxy the median worker's value of non-legally required benefits seems to be a reasonable approximation to employee value. For example, use of this aproximation along with an estimate of after-tax wages to compare the "typical" employee after-tax value of compensation in two industries would appear reasonable. If there are differences in median after-tax wages and noncash benefits between industries, interesting research could be done to determine what the source of the differences is—for example, different median characteristics of the work forces in those industries, differences in median job amenities, or some other disparity. Employer cost as a proxy for how the median employee's value of benefits has changed over time also seems reasonable.[32]

In contrast, the use of employer cost as an approximation in distributional analyses could be highly misleading. Such studies typically focus on the well-being (proxied by some measure of income) of people at varying income levels or family structures. But as we have argued, income and family structure are themselves associated with variations in employee values, and these variations may be considerable. Further, and perhaps more important, the available employer cost measures refer to the "typical" worker in broadly defined industries (nine) and occupations (three). But employers often actually pay more for benefits provided to some types of workers than for others. For example, they may make higher pension contributions for more highly compensated workers or pay a higher cost for family health insurance policies than for single coverage. Thus, imputing these average employer costs to individual observations in a household file to do distributional studies would clearly yield misleading results. There is no theoretical basis for concluding that such comparisons using after-tax cash wages plus the average employer cost for benefits would provide a better proxy for the value of compensation than would use of after-tax cash wages only.[33] Empirical evidence of the extent of the bias involved would, of course, be useful.

More research on employee values is needed. In particular, empirical estimates of the difference between employee value and employer cost for workers of different demographic characteristics would provide evidence for whether the possible discrepancies are significant enough to preclude the use of employer cost as a measure of employee value in distributional studies. Empirical estimates of whether, and how much, employer per-unit cost differs from market price for the median worker in different industries, occupations, areas, and so forth, would provide evidence as to whether cross-sectional analyses using employer cost as the measure of employee value are reasonable. For these purposes, additional data and methodological research are essential. □

Footnotes

[1] The focus is on the employee value of employer-provided noncash benefits, as opposed to job characteristics, such as safety, security, and cleanliness. For the purposes of this article, the chief distinction between the two is that there exists an explicit market for noncash benefits. As a result, while many measurement issues are conceptually the same for both job dimensions, in analyses involving benefits, one must always take into account the fact that the employee could have purchased the benefit on his or her own.

[2] A detailed discussion of the notion of cash-equivalent value and its relation to measurement issues in the case of Government-provided transfers is given in Timothy M. Smeeding, "Alternative Methods for Valuing Selected In-