

ON-THE-JOB LEISURE AND ASYMMETRIC OBSERVED-EFFORT DISTRIBUTIONS*

David L. Dickinson

Dept. of Economics
Appalachian State University

ABSTRACT

When employers observe imperfect measures of worker effort, theorists typically assume that observed effort is unimodal and symmetrically distributed. Though observable effort may be distributed in different ways within a work day, for example, available field data on these effort distributions are rare. The symmetry assumption is largely untestable as a result. This paper presents empirical data from two experimental work environments that question the validity of such assumptions. For these piece-rate work environments the author finds that observed effort is significantly negatively-skew (i.e., modal > mean effort). The author's hypothesis is that on-the-job leisure causes this skewness in observed effort distributions. There are both theoretical and practical implications of this asymmetry. Some implications from the theoretical agency literature that we discuss include: self-selection into rank-order tournaments, optimal wage spreads in rank-order tournaments, and optimal wage contracts with asymmetric information.

*The author thanks Ron Oaxaca, Takao Kato, Jill Tiefenthaler, Jyoti Khanna, Gerald Sohan, and participants in the Economic Science Association meetings for valuable comments on this paper.

1. Introduction

A worker's supply of effort in the work place is of fundamental importance to the success of the firm. Since it is often difficult and/or costly to measure a worker's true effort (henceforth, just "effort"), a variety of incentive schemes rely on the employer's observations of worker output. While rarely addressed in theory and/or empirical work, the consumption of on-the-job leisure (or non work-related activities) is a reality in virtually every work place. A distinctively modern twist to decreased work effort is seen in Gordon (2000), which highlights how the expansion in use of the Internet is affecting labor productivity as workers access entertainment, auction, and financial trading sites, among others. The question is not whether on-the-job leisure occurs, but rather how it affects a worker's observable effort. When a worker consumes on-the-job leisure at points throughout the work day, this leisure then affects the daily effort distribution from which an employer may sample. On-the-job leisure is therefore a realistic example of how worker may choose to distribute work effort throughout the day in a way that is not necessarily random or symmetric.

The theoretical agency literature has addressed this issue of imperfectly observed employee effort, and a typical assumption is that observed effort is symmetrically distributed around an employee's mean effort level for the work period (e.g., a day, month, year, etc.). It is also true that the optimal contract offered to a worker often depends critically on the symmetry assumption. Different versions of random and symmetric production or observed effort distributions can be found in Lazear and Rosen (1981), Malcomson (1984), Parsons (1986), and Yun (1997), among others.¹ Yun notes that the assumption of a well-behaved error term is

¹ Lazear and Rosen assume that lifetime output is based on both effort and a well-behaved random error term, Malcomson relates output to effort with a symmetric and unimodal error term, Parsons gives several models that

“...standard but critical.” It is clear that, when using a stochastic production technology to link effort and output, the standard theoretical practice is to incorporate well-behaved error terms that generate symmetric distributions.²

In this paper, I raise the possibility that distributions of observable worker effort may be skewed, and the hypothesis is that on-the-job leisure is the cause. An *ex ante* assumption of this paper is that workers consume on-the-job leisure.³ Not only is this hypothesis supported by anecdotal evidence, but I also present empirical evidence from two controlled laboratory work environments in which on-the-job leisure is precisely measured. In this case, the experimental environment generates empirical data that is more detailed than typical field data, such as data on intra-period work effort. Typical proxies for work effort in field data are more aggregate responses (i.e., responses per labor period) that do not allow for distributions of intra-period effort to be studied. Surveys are also subject to possible recall or self-reporting bias.⁴ Laboratory methods can limit these biases in the data and often generate data very specific to the research at hand.

The implications of asymmetric effort distributions, such as those found in the lab environment, are discussed more thoroughly in Section 4 of the paper. Included among these are the implications for several results from the agency theory literature: self-selection into rank-order tournaments, optimal wage spreads in rank-order tournaments, and optimal wage contracts

assume the value of a worker’s output is a function of effort and a random zero-mean error-term, and Yun includes a well-behaved error term in relating an employee’s true effort to the employer’s observation of that effort.

² An exception brought to my attention by Ron Oaxaca is that of frontier regression techniques applied to production technologies. Here, a truncated error distribution may reflect the assumption that output levels cannot lie above the frontier.

³ Dickinson (1999), for example, uses a simple theoretical framework to generate an optimal implicit on-the-job leisure function.

⁴ Greenberg, et al. (1981) uses hours of work as a proxy for work effort in the Negative Income Tax (NIT) experiments and concludes that there is “substantial” labor supply underreporting since the experimental payments were, in part, determined by this variable.

with asymmetric information. This is not an exhaustive list of the implications of an asymmetric effort distribution, but it should suffice to highlight the point that optimal contracts derived from such theories (and workplace policies) can be sensitive to the form of the effort distribution. The objective of this paper is to reveal the theoretical *and* practical importance of a seemingly innocuous theoretical assumption that can be violated under certain plausible conditions.

2. Effort and on-the-job leisure

Suppose that worker output (or production rate), y , is the worker's observable effort to the employer. The theoretical agency literature recognizes that worker effort is not perfectly observable to the employer. Let us model this by simply assuming $y=f(e)$, where e is maximum effort and $f(e)$ is a random variable with an associated p.d.f. possessing certain properties. This makes explicit the often *implicit* assumption that true effort on the job is fixed at its maximum level—this is true in models where hours of work is the only choice variable. Worker output, y , is therefore a random variable. A standard assumption would be that y is a symmetrically distributed random variable (e.g., stochastic due to production technology shocks). On-the-job leisure, l_w , during the work period (which can be modeled as leisure per day, week, year, etc) can be thought of in at least two different ways. First, if l_w is consumed at a *constant rate* throughout the work period, then l_w can be thought to parametrically shift the distribution of worker output. Under this assumption l_w would not alter the shape of the output distribution. If, on the other hand, on-the-job leisure is not consumed at a constant rate throughout the day, then let $g(l_w)$ be a random variable, and its associated p.d.f. describes how on-the-job leisure is distributed throughout the work period. Let us now model observable worker output as a function of maximum effort *net* of on-the-job leisure, or $y=f(e)-g(l_w)$, which is a random variable.

If one assumes that both $f(e)$ and $g(l_w)$ are distributed normally, then y would be a symmetric, normal distribution of worker output. Under such assumptions, a stochastic production technology maps *net* effort to output, but the shocks to net effort (maximum effort minus on-the-job leisure) are still random and symmetric as is typically assumed. On-the-job leisure does not affect maximum possible work effort, but it does affect net effort and observable worker output within this framework. Peculiarities in on-the-job leisure can both parametrically shift the observable effort distribution by the mean level of $g(l_w)$ as well as shape the observable effort distribution via the shape of the $g(l_w)$ distribution. This explicit (versus implicit) recognition of on-the-job leisure does not affect any of the points to be made in this paper, but the distinction is explicit to facilitate discussion of output as a function of maximum effort *net* of on-the-job leisure.

A key question logically follows: How might on-the-job leisure be distributed throughout the work period? While a symmetric distribution may appear a reasonable assumption, I hypothesize that a skewed distribution of the p.d.f. associated with $g(l_w)$ is also a reasonable assumption. Consider how l_w might be consumed. Is it through mini-breaks during the work period, conversations with one's colleagues, or a nap at the office? Do individuals tend to sprinkle on-the-job leisure evenly throughout the day, or do they hoard it to be consumed in occasionally large quantities? Put simply, if individuals consume l_w in large quantities at times, then this will positively skew the $g(l_w)$ distribution, much like income distributions are positively-skew due to outlier high-income individuals. For example, $g(l_w)$ might be a Poisson distribution. A Poisson distribution would reflect the fact that $g(l_w)$ is truncated at zero, but not truncated on the right-tail of the distribution—a reasonable assumption is that on-the-job leisure is bound only on the lower end by zero. By assuming a positively-skew $g(l_w)$ distribution, the

resultant distribution of observable effort, y , would then be negatively-skew (modal $>$ mean y) if $f(e)$ is well-behaved. The basic conjecture in this paper is that the distribution of observable worker effort (output) may be negatively-skew in the presence of certain patterns of on-the-job leisure consumption.⁵

Standard measures of on-the-job leisure and distributions of worker output are difficult to explore with existing field data. The next section presents some empirical distributions of per-period production rates from two laboratory work environments. At the very least, such experiments provide data not otherwise available, and the data may help guide how we think about observable effort in naturally-occurring work environments. Further, since the data were originally generated to test other hypotheses, as Roth et al. (1988) note, “..there is a sense in which it [such data] is not all fully ‘experimental data,’ ...” (p. 808).

3. Evidence for the Asymmetry of Observed Output

The Data Set

The data set used in this article is presented first in Dickinson (1999)⁶, where subjects complete a task repetitively for piece-rate wages. Each subject types paragraphs, subject to a quality control check, and receives on average \$.30 per paragraph that takes the average subject

⁵ This is not the only thing that could skew the distribution of observable output. The production function itself might skew the distribution of effort if technology, for example, asymmetrically shocks worker effort within a work period. On-the-job learning could also skew the output distribution, although one would assume that an employer would choose a long enough work period (or discard any output observations from the initial “learning” period) so that learning would not significantly affect any incentive contracts. Fatigue could also skew the observable effort distribution. The results presented in this paper have been examined to ensure that neither fatigue nor learning is the sole cause of the skewness observed in the empirical evidence (evidence available on request), though such items could also be important in naturally occurring work environments. Finally, if worker expectations or aspirations change within the work period, then this could also affect the shape of the effort distribution.

⁶ The principle aim of this 1999 study is to decompose a wage change into substitution and income effects and to analyze the resulting changes in daily work effort and/or time spent working.

about 5 minutes to type.⁷ Subjects also received a fixed nonwage payment for each day's work, and they each participated in four experimental workdays within a ten-day period. Two data sets are discussed and presented. Depending on which experimental design is employed, subjects either work for a fixed 2-hour workday (the Intensity Experiment) or they were allowed to leave at their choosing (the Combined Experiment). The Intensity Experiment only allows choice of work intensity, whereas the Combined Experiment allows choice of both work effort and hours of work (i.e., on-the-job and *off-the-job* leisure). The piece-rate wage structure of the work environments may limit the generality of this data set, but piece-rate wages (or quasi piece-rate) are still found in a variety of settings (e.g., contract work, fruit and vegetable harvesting, jobs where a substantial portion of pay is in the form of tips, and a wide variety of sales jobs that pay on commission).

Observed effort within the workday is precisely measured in both experiments. A subject, after typing his paragraph on the computer terminal, prints the paragraph to a central printer and the printer places a time stamp on the paragraph. The experimenter checks the paragraph for errors (up to five were allowed without penalty) and records the time stamp as an exact measure of how long it took to complete that unit of output. The result is a series of paragraph production times for a total of 41 subjects (26 in the Combined Experiment and 15 in the Intensity Experiment) for each of 4 days.⁸ The computerized collection of the production time data is not only precise, but it is void of self-reporting, bad memory, or recording bias. Since all experimental subjects work independent of each other and were not allowed to communicate with other subjects, there is no possibility that the data includes worker attempts to

⁷ If one were to ignore the fixed nonwage payment, the average subject earned less than \$4 per hour of incentive pay. This was enough to motivate subjects, however, as the effort response to a wage change is statistically significant.

sabotage others' output. Finally, since subjects were unaware of whether their work days would include a wage increase or decrease, there was no incentive to strategically alter daily effort in order to somehow maximize *total* experimental earnings.

Different subjects possess different typing abilities as one might expect, and so the production data is not immediately comparable across individuals. To pool the data, I first convert the data to production rates, y/t —the number of paragraphs, y , produced per $t=5$ minute period of time—and then standardize (i.e., subtract the mean and divide by the standard deviation) each day's distribution of production rates for each individual.⁹ Production rates are therefore an empirical measure of subject effort. The choice of 5-minute time period to define production rates is arbitrary, but the conversion of production times to production rates is necessary to yield an observable *output* variable as opposed to the paragraph production times. It should also be noted that, by creating a standardized distribution of production rates per individual *per day*, the implicit assumption is that the subjects' choice period is the 2-hour workday. This is reasonable given that the purpose of the experiments was to alter wage and/or nonwage income across (not within) different workdays. Pooling the data across individuals, there are a total of 1493 observations of production rates in the Intensity Experiment and 1344 in the Combined Experiment.

In sum, this data set seems well-suited to the question at hand. Subjects' observable output is clearly a function of their work effort and on-the-job leisure choice. While fatigue and/or learning might also affect observable effort distributions, these experiments control for learning across experimental days (see Dickinson (1999)) and do not exhibit signs of any

⁸ Though Dickinson (1999) analyzes only the data from the final 3 work days in the experiment, the initial work day is also included in the pooled data analysis of this paper since all data is normalized.

⁹ This is similar to assuming that an employer possesses some individual-specific ability information and takes this into account in incentive contracts (e.g., incentives handicapped for different ability levels).

average upward or downward trend in the production rate data—signs which might signify learning or fatigue is manifest in the data.¹⁰ Though mean production rates often differ across experimental days and as wage incentives changed and across subjects, the data can be pooled to examine output *distributions* given our standardization of the data.

Data analysis and the on-the-job leisure hypothesis

What is clear from Figures 1a and 1b is that the distributions of production rates (i.e., empirical effort) for each experiment are *negatively skew*. When testing for symmetry as well as for normality of the distributions, one can reject the null hypothesis that the distributions in Figures 1a and 1b are normal (Kolmogorov-Smirnov test, $p < .01$ in both cases).¹¹ In testing for symmetry of the distributions, we reject the null hypothesis of symmetry (signed rank test, $p = .04$) for the Intensity Experiment data in Figure 1b, but we marginally fail to reject the null hypothesis of symmetry ($p = .12$) for the Combined Experiment data in Figure 1a.¹² This result is actually quite intuitive when one recalls that subjects can leave the experiment early in the Combined Experiment (i.e., they have an *off-the-job* leisure option). In the Intensity Experiment, subjects were required to stay the full 2-hours. If one assumes that off-the-job leisure is preferred over on-the-job leisure—this seems reasonable given the additional leisure options available *off-the-job*—then when work hours is a choice variable, workers prefer to consume off-the-job leisure rather than on-the-job leisure. Many observations of per unit effort less than average effort (i.e., more on-the-job leisure) would intuitively occur more often when workers

¹⁰ Additional data evidence is available on request.

¹¹ The same results are found using a variety of other test statistics for normality, such as the Shapiro-Wilk, Cramer-von Mises, and Anderson-Darling test statistics.

¹² I also perform a χ^2 test of the null hypothesis that the two empirical distributions come from the same population, and the result is to reject the null hypothesis ($p < .01$). In using this test, some of the outlier data categories are aggregated so that no empty categories exist.

have fixed-hours schedules. On-the-job leisure is therefore a more likely concern when workers are constrained in their hours of work choice.

The data presented in Figures 1a and 1b offer evidence in support of skewed effort of the subject population, but the theoretical literature often makes assumptions about individual effort. In response to any concerns that the data are skewed by some individuals, the data were also examined at the individual level. The full-distribution Kolmogorov test is used to test each individual production day cumulative distribution function (CDF) against two hypothesized distributions: the standard normal, and a low-variance symmetric distribution with its mass highly concentrated about the distribution's mean. The empirical effort (production rate) CDF will lie mostly below, or to the right of, the hypothesized symmetric CDF if individual effort is negatively-skewed similar to Figures 1. The Kolmogorov test utilizes the greatest vertical distance between the CDFs as the test statistic, and a large distance leads to rejection of the null hypothesis that the two CDFs are equivalent.

In comparing the empirical CDF and the standard normal CDF, the greatest difference between the CDFs is when the empirical distribution lies below the standard normal CDF in 69 of 104 cases (66% of the time) for the Combined Experiment (i.e., negatively-skewed as is the population CDF). The difference is statistically significant at the 10% level in 22 of the 69 cases (32%). For the Intensity Experiment, negative-skewness occurs in 50 of 60 cases (83%), but is significant in only 10 of the 50 instances (20%). So, while negative-skewness in the empirical production rate data dominates at the individual level, it is significant in a minority of those cases. However, the standard normal distribution is only one example of a symmetric distribution, and we could also test the individual-level data against a lower variance test distribution to highlight whether there is skewness in the data. Compared to the standard normal

distribution, our low variance test distribution has half its mass located just below and half just above the distribution mean.¹³ Utilizing this hypothesized distribution, we now see the negative-skewness in 86% of the Combined Experiment data (significant skewness 61% of the time) and 100% of the Intensity Experiment data (significant skewness 80% of the time). So, it is clear that significance of the skewness is a function of the test distribution hypothesized, but overall the evidence is clearly in favor of skewness at the individual level similar to that in Figures 1.

A central conclusion in Dickinson (1999) is that on-the-job leisure is enjoyed in both experimental designs. In addition to the evidence in Figures 1, I support the claim that on-the-job leisure is not consumed in a random and symmetric fashion throughout the workday by highlighting a couple of salient, albeit atypical, subjects from the Intensity Experiment. For Subject 2 on Day 3, the mean paragraph production time is 6.7 minutes, with an outlier of 36.7 minutes. While such an outlier does contribute to a high variance in that day's production time distribution, that particular data point is still far into the left tail of the production rate distribution.¹⁴ An equivalent production time in the far right tail of the distribution would have to be negative, which is an obvious impossibility. Similarly, Subject 4 in the Intensity Experiment had a mean production time of 7.2 minutes on Day 4, and an outlier of **59 minutes**.¹⁵ Indeed, both of these subjects were witnessed falling asleep on their respective keyboards as a manifestation of their choice to hoard their daily on-the-job leisure.¹⁶

¹³ Compared to Figures 1a and 1b, the low-variance test distribution has half of its mass in each of the two bins on either side of the distribution mean.

¹⁴ Omitting these extreme "outliers" does not significantly alter the overall shape of the production rate distribution shown in Figure 1b.

¹⁵ One could only consume on-the-job leisure and no off-the-job leisure in the Intensity Experiment since the workday was fixed at 2 hours. Nonetheless, consumption of on-the-job leisure was still present in the Combined Experiment and, while outlier production times were not as large in absolute terms, the production time distributions were still skewed, just not statistically significantly so as they were in the Intensity Experiment (see Figures 1).

¹⁶ I explored the possibility that intra-day learning of the typing task was responsible for the skewness of the production rate distribution. A separate examination of the data from the initial 30 minutes of the work day reveals that both the initial work day and the remaining work day generate skewed data. As such, even if skewness in the

The external validity of laboratory data is always a concern. The data from these experiments do, however, offer the unique opportunity to examine production rates of individuals who face real incentives in choosing between labor and leisure. In work environments where employers are assumed to imperfectly observe worker effort and sample from a distribution of observable output, these data imply that work place policies and optimal contracts may be affected by the presence of on-the-job leisure.¹⁷ The data also highlight that concerns over the effects of on-the-job leisure are perhaps most warranted in workplaces where hours of work are *not* at the discretion of the worker, which is likely the case in many work environments.

4. Implications

In this section I highlight several implications of an asymmetric observed effort distribution. Though asymmetric distributions of observed effort likely have implications for most theories that assume a symmetric distribution, by discussing at least a few implications in greater detail I hope to highlight the importance that the asymmetry may have in the workplace for both theorists and practitioners.

Some conclusions can be readily seen even in the absence of an explicit theoretical model. For example, if an employer incorrectly assumes that per-period output is symmetrically distributed and bases reward and/or sanction policies on sampled production rates that are above or below some threshold, then reward and penalty probabilities will be miscalculated. One way

initial 30 minutes is due entirely to learning (which may or may not be the case), on-the-job leisure appears to still play a significant role in the data for the remaining part of the work day.

¹⁷ The data used in this analysis may not be independent observations. A nonparametric one-sample runs (chi-squared) test on each separate experimental day reveals that about 10% of the work days in the Combined Experiment and 25% of the work days in the Intensity Experiment display nonrandom runs of production times either above or below the mean production time for that work day. In the minority of the cases when workday data are not random, subjects group together several observations of above- and below-average effort during the work day.

to avoid this problem would be the practice of employers to promote, for example, the top 5% of employees (ranked by observed effort). Greater information (i.e., more frequent sampling) can limit such miscalculation, but gathering such information is costly. This highlights that policies based on absolute effort thresholds will not generate the *ex ante* desired outcomes (e.g., number of dismissals or bonuses) if symmetric output rates are mistakenly assumed. Other workplace implications can be derived from specific theoretical models, but even the relatively informal discussion that follows will highlight that even optimal contracts based on rank orderings of individuals are not immune to the effects of an asymmetric output distribution.

Malcomson (1984) develops a theory of contracts and effort in which observed (symmetrically distributed) effort of an employee determines whether or not that employee is promoted to earn a higher wage in period two of a two-period contract. This choice of the firm's optimal contract is shown to depend upon, among other things, how observed effort responds to the probability of promotion (p.497, equation (10)). For the assumed symmetric and unimodally distributed observed effort, Malcomson shows that when observed effort is above average, *true* effort (or, "net" effort in our terminology) increases in the probability of promotion, P . His conclusion, however, is really that effort is increasing in P so long as the observed effort p.d.f. declines, which is only for effort *greater* than the mean for a symmetric unimodal distribution.

Figures 1a and 1b show that when observed effort is negatively-skew, there exists levels of observed effort greater than the mean level such that optimally chosen true effort would be declining in P (since the p.d.f. is increasing) according to Malcomson's model. The implication for the optimal contract that the firm would offer is as follows: Consider that firms "make up" the higher promotion wage through higher worker effort, and firms must set P in the employment contract to maximize expected utility. Asymmetry in the observed effort distribution creates a

range of observed effort for which the firm should decrease rather than increase P in designing its optimal contract.

As a second example, consider Lazear and Rosen (1981)—a seminal paper on rank-order tournaments. Per period output is a function of average output (or effort) and a random well-behaved error term, ε . Per period output is therefore analogous to observed effort.¹⁸ A rank-order tournament is shown to possess the same efficiency properties as piece-rate wages for risk-neutral workers. Wages offered in the tournament are either W_1 or W_2 with $W_1 > W_2$, so that W_1 is the fixed prize going to the tournament winner. A key result is that there is a unique equilibrium spread ($W_1 - W_2$) that maximizes workers' expected utility. It is shown that for a risk-neutral firm, when ε is normal, the optimal prize spread varies directly with the variance of the output distribution. This result relies on the fact that the difference in workers' levels of investment (e.g., true effort) is compared to the difference in the workers' error terms (i.e., the random shock to each production technology—their equation (3) on p. 845). Normally distributed error terms guarantee that this difference $\varepsilon_i - \varepsilon_j = \xi$ is also normally distributed.

If output of each worker is determined by a stochastic production technology that asymmetrically shocks effort (e.g., because of on-the-job leisure), this difference in errors is not likely to be symmetric—its precise form, or even whether it is positively- or negatively-skew, cannot be determined without further assumptions on the nature of the shocks ε_i and ε_j .

However, their 1981 result is somewhat more general than what they show. If one considers a positively-skew gamma distribution (as the density function of ξ that results from the difference

¹⁸ In this paper lifetime output is assumed to depend on skill or average output (e.g., effort) as well as lifetime luck represented by a well-behaved error term. Lifetime output is considered to reduce the problem to a single period to avoid sequencing effects. Of course, there is no problem in assuming that lifetime effort could also be asymmetrically distributed (even with a symmetric “luck” term). Though not a lifetime in any real sense, the experimental data is still asymmetrically distributed at our level of aggregation across work days.

in asymmetrically distributed errors, $\varepsilon_1 - \varepsilon_2$), then it is still the case that the optimal prize spread varies directly with the distribution of the error terms, and hence the variance of ξ , $E(\xi^2)$.¹⁹ This result notwithstanding, the amount by which the risk-neutral firm would adjust the optimal prize spread is affected by the asymmetry of ξ . For example, assume that employers incorrectly perceive mean effort to be modal effort, and the density function of ξ is, in reality, less sensitive to a change in $E(\xi^2)$ when evaluated at the mean than at the mode (as is true for many parameterizations of the gamma distribution, e.g.). An increase in $E(\xi^2)$ will then induce the employer to increase the prize spread by too much, thus resulting in a *non-optimal* prize spread.

As a final example, consider the issue of efficient self-selection into simple rank-order contests (SRC's) as set forth by Yun (1997). As employers set an observable effort threshold below which workers are penalized, there exists a given probability of being penalized, and this probability decreases in worker effort. A salient result in Yun is that penalty probability functions must be convex for the existence of efficient self-selection of individuals of differing abilities into tournaments of their own type (high or low ability workers). This convexity implies that when "average" effort increases above the penalty threshold, the probability of being penalized decreases at a decreasing rate. Imagine a given penalty threshold, and a worker increases maximum effort in a way that parametrically shifts the observable output distribution—this would be possible with additional technology or human capital investment by the worker. With a typically assumed unimodal and symmetric distribution of observable effort, the penalty probability function is convex for all true (net) effort levels above the penalty

¹⁹ The proof of this is available from the author upon request.

threshold. Under the empirical negatively-skew effort distribution, the convexity region extends to lower effort levels.²⁰

Figure 2 shows the difference in the penalty probability functions assuming the two different observed effort distributions. This result implies one of two things. The first possibility is that the existence of a first-best SRC holds at lower net effort levels than previously thought. The second is that for the existence of a first-best SRC at the effort levels shown by Yun, a *higher* proportion of workers could be penalized (or, alternatively, a higher penalty threshold established) under the asymmetric output distribution than under the symmetric one.

Strictly speaking, the evidence in this paper does not imply that tournament or other settings will generate asymmetric observed effort distributions since the data presented are from piece-rate settings.²¹ The point of utilizing varied examples is to stimulate thought on potential implications. Theory always establishes assumptions prior to developing implications. A theory or policy that performs well is one whose results are robust with respect to the relaxation of assumptions. As shown in the above examples, the assumption of symmetric observed effort distributions is often a key assumption, and to relax this assumption implies that theory and/or workplace policy results would often change.

5. Concluding Remarks

Existing studies that hypothesize a distribution of worker's observable effort typically assume a symmetric distribution, such as the normal distribution. Theoretical results often hinge on this assumption, as would workplace policies based upon such theories. On the other hand, if

²⁰ In Yun, the different choices of effort that generate the desired penalty probability functions are true effort choices—the mean of the effort distribution in Yun. In this present article's framework two items can change: either maximum effort potential (e.g., through use of new technology) or the chosen mean level of on-the-job leisure. Either of these would parametrically shift the observable output distribution. The point is that the shifting of asymmetric versus symmetric distributions does matter in terms of the theoretical implication.

²¹ The author thanks an anonymous reviewer for highlighting this point.

observable effort or output follows an asymmetric distribution then this should be of interest to both theorists and practitioners. Such an asymmetry in the per-period distribution of worker output can result from on-the-job leisure. On-the-job leisure can be manifested by periods of low output rates that are not matched by equally-sized bursts of high output rates, and this extends the left-tail of the observed effort distribution.

Dickinson (1999) provides a controlled set of experiments that generate data on observable production rates within the workers' choice period. The data reveal an asymmetric distribution of observable effort in the work environments. The asymmetry is statistically significant in a work environment where subjects are constrained in their hours of work choice. This implies that, for those environments where on-the-job leisure is most likely to be consumed, agency theory and workplace policy should be concerned about the implications of the observable effort asymmetry. These experimental work environments represent piece-rate wage environments, and some would argue that high or low effort is of no real consequence to the employers of piece-rate workers since the piece-rate is itself the work incentive. However, piece-rate employers still have an interest in labor productivity since labor inputs are often combined with other complementary inputs in the production process. Further empirical work is still advised to establish the generality of asymmetric observed effort distribution in many different work environment (e.g., hourly, tournament, hybrid environments, etc).

The implications of incorrectly assuming a symmetric distribution of observed effort can be found in many theories or policies based upon the assumption. These limitations range from theoretical to practical and, while not an exhaustive list, the following have been highlighted: In setting wage contracts and promotion probabilities (P), asymmetric observed effort affects whether or not employee effort is increasing or decreasing in P —a point which is crucial to

choosing the optimal wage contract. It is also shown that in rank-order contests, an asymmetric observed effort distribution can imply an equilibrium prize spread that is less sensitive to the variance of observed effort than previously thought. Finally, where minimum effort requirements are utilized, the existence of efficient worker selection into simple-rank contests occurs with higher (than previously thought) penalty probabilities (and, hence, higher minimum effort requirements). Many employers also offer positive incentives to workers surpassing some level of effort. To the extent that employers have an *a priori* idea of reward and/or penalty probabilities for their workplace policies, these probabilities will be biased in predictable ways when output is asymmetrically distributed. Employers may penalize too few workers or reward too many or set unattainable goals for employees, all of which can be costly miscalculations.

The goal of this paper is not to discredit agency theory, but rather to shed some light on the possibility and implications of asymmetric observed effort distributions. These can occur under very reasonable workplace assumptions about on-the-job leisure, and negatively-skew output rate distributions are documented in the laboratory work environments discussed in this paper. Future research should perhaps examine a wider variety of experimental work environments as well as attempt to gather richer field data sets that could provide intra-period output data in an attempt to establish the generality of asymmetric observed effort distributions.

References

- Dickinson, D.L. (1999). "An Experimental examination of labor supply and work intensity." *Journal of Labor Economics*, 17(4), part 1: 638-670.
- Gordon, R. J. (2000). "Does the 'New economy' measure up to the great inventions of the past?" *Journal of Economic Perspectives*, 14(4): 49-74.
- Greenberg, D., Moffitt, R., and Friedmann, J. (1981). "Underreporting and experimental effects on work effort: Evidence from the Gary Income Maintenance Experiment." *Review of Economics and Statistics*, 63(4): 581-589.
- Lazear, E. and Rosen, S. (1981). "Rank-order tournaments as optimum labor contracts." *Journal of Political Economy*, 89: 841-864.
- Malcomson, J. (1984). "Work incentives, hierarchy, and internal labor markets." *Journal of Political Economy*, 92: 486-507.
- Parsons, D. O. (1986). "The Employment relationship: Job attachment, work effort, and the nature of contract." In Ashenfelter, O. and Layard, R. eds. *The Handbook of Labor Economics*, volume 2. Elsevier Science Publishers B.V.
- Roth, A., Murnighan, J.K., and F. Schoumaker (1988) "The deadline effect in bargaining: Some experimental evidence." *American Economic Review*, 78(4): 806-823.
- Yun, J. (1997). "On the efficiency of the rank-order contract under moral hazard and adverse selection." *Journal of Labor Economics*, 15(3), Part 1: 466-494.

FIGURE 1a

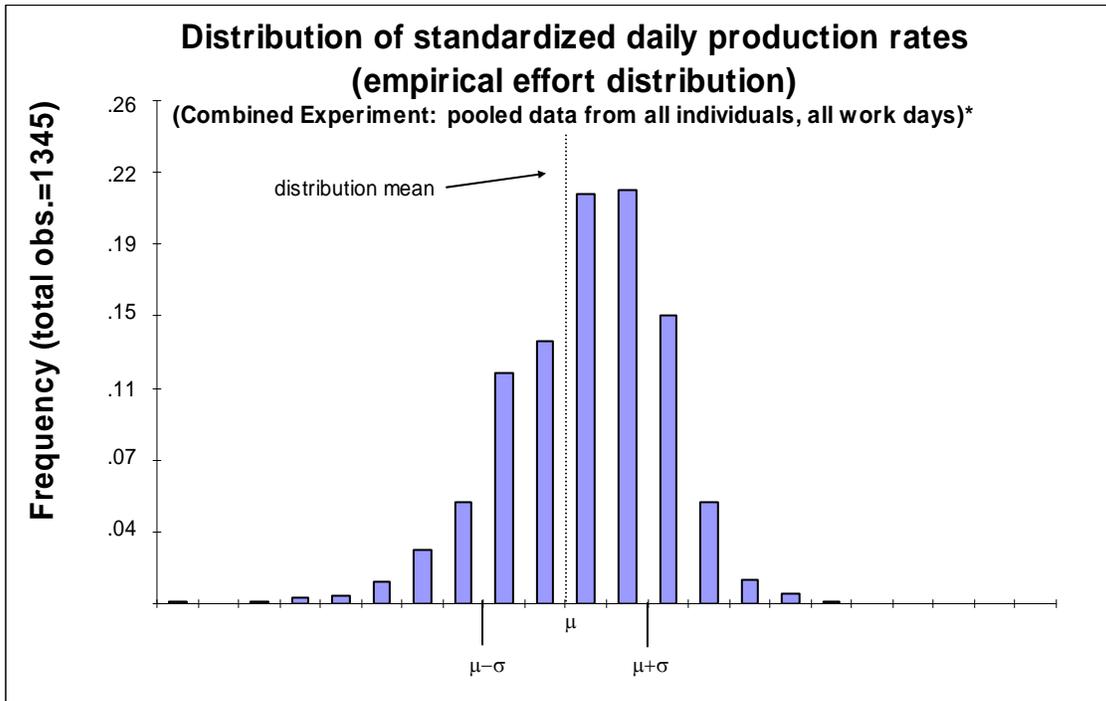
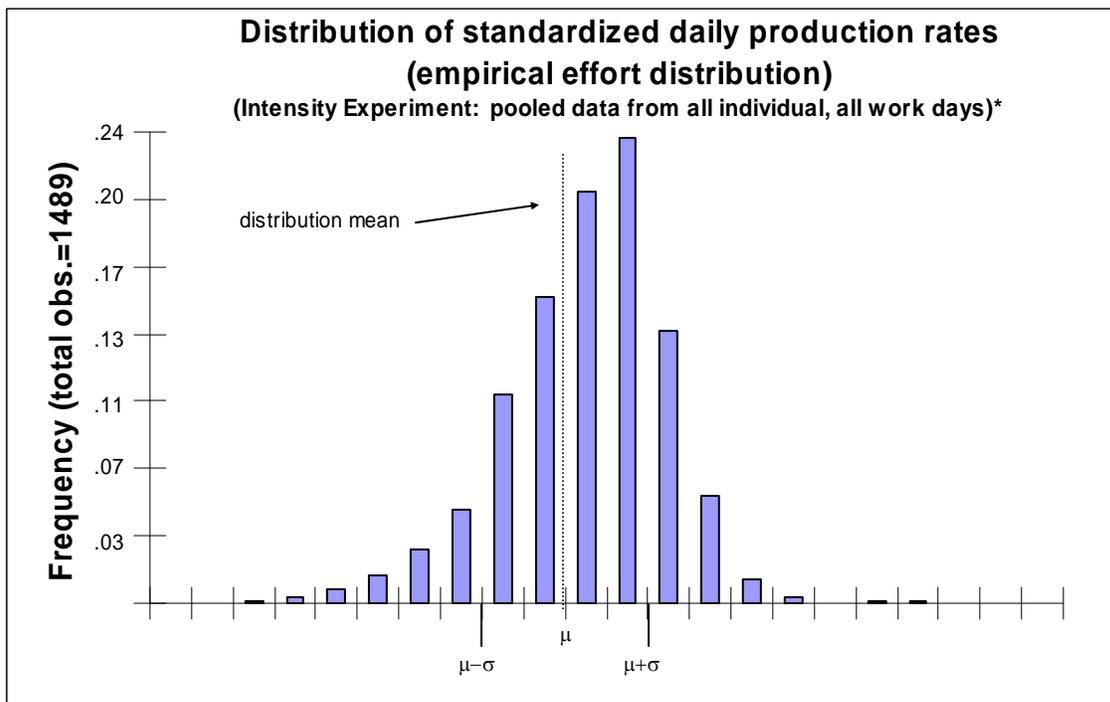


FIGURE 1b



*Referenced experiments are from Dickinson (1999)

FIGURE 2

