Measuring Temporary Labor Outsourcing in U.S. Manufacturing*

Marcello Estevão
Board of Governors of the Federal Reserve System

Saul Lach
The Hebrew University of Jerusalem and NBER

May 8, 2000

Abstract

Temporary help supply (THS) employment has increased dramatically in the last ten years. However, there is only scant evidence on the industries that are hiring THS workers. Without this information, it is difficult to determine the reasons for the surge in temporary help usage in a particular sector. In addition, not accounting for temporary workers may result in biased interpretations of the cyclical and long-term properties of sectoral employment. These issues are particularly relevant to the manufacturing sector since anecdotal evidence suggests that manufacturers have substantially stepped up their demand for THS workers since the mid-1980s. Here we provide estimates of the number of THS employees in manufacturing from 1972 to 1997. Our estimates are based on a new methodology to put bounds on the probability that a manufacturing worker is a THS employee. We verify that manufacturers have been using THS workers more intensively during the 1990s. This new trend partly explains the sluggish recovery of non-THS manufacturing employment in the 1990s. Finally, not accounting for THS hours overstated the increase in average annual manufacturing labor productivity by 0.30 percentage point during the 1991-1997 period.

JEL codes: E2, J2, D2, C14

*This is a revised version of the NBER working paper 7421. Contact author: Saul Lach, Department of Economics, The Hebrew University, Jerusalem 91905, Israel. We thank Chuck Manski for helpful discussions at an early stage of this project, Cynthia Bansak, Carol Corrado, Anne Polivka, Dan Sullivan, Beth Anne Wilson and Joyce Zickler for their comments and Nicole Nestoriak for excellent research assistance. We also got valuable suggestions from seminar participants at the 1998 NBER Summer Institute, the 1998 AEA meetings in Chicago, IL, the Federal Reserve Board, the Hebrew University, Tel-Aviv University, and the University of Cyprus. Saul Lach acknowledges financial support from the Employment Policies Institute. This paper presents the authors' own views and not those of the Federal Reserve System or of any members of its staff.

1 Introduction

The recent debate on the extent of labor outsourcing by manufacturing firms has been hampered by the absence of good data. In this paper we estimate the number of temporary workers in manufacturing firms that are not hired directly by the firm but instead work under contract with firms that are primarily engaged in supplying temporary help to other businesses. More specifically, we estimate the number of "temporary help supply" (THS) employees–individuals on the payrolls of service sector firms–working in manufacturing. We construct annual estimates from 1972 to 1997 and use them to trace the evolution of this form of labor outsourcing. We also use estimates of the number of THS hours in manufacturing to correct the official measures of manufacturing labor productivity.

The hiring of THS workers is one aspect of the general trend toward flexible, market-mediated, work arrangements by firms. Tasks that formerly were performed by workers hired directly by the firm are now done under contract with firms in the business service sector. Such arrangements include outsourcing of various support services (e.g., computer maintenance, accounting, etc.), subcontracting specific tasks in the production process, and using temporary employees. Many reasons have been advanced to explain the rapid spread of THS arrangements. These include the potential for employers to implement a new lower-wage rate in a two-tier wage structure by contracting with intermediaries that pay less for similar work, to realize scale economies due to specialization in the provision of specific tasks, to increase productivity given that THS workers may be better screened or trained than temporary workers hired directly by the firm (Autor, 1998; Polivka, 1996), and the potential to facilitate more rapid changes in firms' level of employment in response to temporary and/or unpredictable changes in demand (Abraham and Taylor, 1996; Golden, 1996).

The increased use of THS workers is evident in the payroll data published by the Bureau of Labor Statistics (BLS). In the last decade, employment in the temporary help supply industry has more than tripled in the United States.[1] Although employment in the THS industry represented only about 2 percent of total nonfarm employment in 1997, it accounted for 10 percent of the net increase in nonfarm employment between 1991 and 1997. Since 1972, employment in the THS industry has risen at an annual

[1] The use of temporary workers also grew rapidly in most OECD countries (International Herald Tribune, September 2, 1997).

rate of more than 11 percent while total nonfarm employment has expanded only at an annual rate of 2 percent (figure 1).

In addition, THS jobs are highly cyclical. While annual changes in nonfarm employment have fluctuated between -1.8 and +5.1 percent since the 1970s, employment changes in the THS industry have ranged from -25 to +32 percent. Moreover, economists have found THS employment to be a leading indicator of overall employment conditions (Segal and Sullivan, 1995).

But, how does the surge in THS jobs affect employment in other sectors? Because the BLS classifies employees by the industry where they are employed rather than by the industry where they are working, previous studies based their answer on very strong assumptions regarding the distribution of these workers across industries.

This paper addresses this problem by developing a methodology for estimating the number of THS workers by industry. We derive non-parametric *bounds* for the proportion of THS employees working in manufacturing. The theoretical approach uses minimal assumptions while exploiting the richness of the Current Population Survey tapes and the Contingent Worker Supplements. The procedure can be readily applied to any year for which data on characteristics of THS individuals are available and, in principle, it can be applied to any level of aggregation. More generally, this technique can also be used in to studies aiming to identify conditional probabilities from observed marginal probabilities.

In our particular application, we find that the bounds on the probability that a manufacturing worker is a THS employee are non-trivial and trend upward. Using the mid-point of the interval generated by the bounds and direct survey evidence for 1995 and 1997 we derive some of the quantitative implications of our estimates.

A simple analysis of reported manufacturing payroll employment data suggests that the period between 1991 and 1997 generated only about 271,000 manufacturing jobs. The inclusion of THS workers more than doubles this figure to 581,000. Manufacturers are estimated to have employed 657,000 THS workers in 1997. Moreover, the decline in manufacturing hours between the local peak in 1989 and 1997–about 1.25 percent–disappears once THS workers are taken into account. We also show that the year-to-year variation in manufacturing THS employment and hours is of an order of magnitude larger than for manufacturing non-THS employment and hours.

Our point estimates can be used to assess the magnitude of the upward bias in manufacturing labor productivity caused by the omission of THS hours from the official payroll statistics. After adjusting for THS hours, labor productivity in manufacturing at the end of 1997 was about 2.5 percent lower than the "reported" level. This correction, while noticeable, explains only a small part of the observed gap between trend labor productivity growth in manufacturing and in nonfarm nonmanufacturing industries.

The paper is organized as follows. After a brief description of the THS industry we formally define the measurement problem in section 3. Section 4 presents our solution to this problem–estimating bounds for the proportion of THS workers–which constitutes the main methodological contribution of the paper. Section 5 presents the data and discusses some issues relevant to the estimation of the bounds, while section 6 presents the estimated bounds on the proportion of THS workers and hours in manufacturing during 1972-1997. Section 7 checks the sensitivity of these estimates to alternative estimation procedures and section 8 derives the quantitative implications summarized above. Brief conclusions close the paper.

2 The Temporary Help Supply Industry

What exactly do firms in the THS industry do? Firms in the THS industry are essentially offering a "business service": they recruit and screen candidates for limited-term jobs, write the contracts, and assume the legal responsibilities of hiring and firing. Most importantly, they administer the payroll of their employees even though the employees are obviously under the direct or general supervision of the business to which the help is furnished.

What types of workers are used by these firms? The March demographic files of the Current Population Survey (CPS) shed some light on the characteristics of workers in the personnel supply services (PSS) industry (SIC 736), the industry that contains the THS industry (SIC 7363). The notion that a "temp" is a woman working part-time in a clerical position, with little job security and lower-than-average wages and benefits, does not seem to be as accurate in the 1990s as it was in the 1970s.

Segal and Sullivan (1995, 1997) report an increase in the proportion of men and blue-collar workers at the end of the 1980s and beginning of the 1990s. Estevão and Lach (1999) showed that this trend has continued well into the 1990s.[2] Segal and Sullivan also reported that while PSS workers are less attached to the labor force than other

[2] In 1992-97, men made up 38 percent of all workers in the PSS industry, but only 28 percent in 1977-87. The fraction of PSS employees working in blue-collar occupations rose from about 10 percent in 1977-87 to about 23 percent in 1992-97.

workers are, a large fraction of them shift into permanent jobs within a year. Finally, although PSS workers on average earn lower wages than workers with similar demographic and educational characteristics do, this wage differential varies widely by occupational group, being largest for blue-collar workers and almost non-existent for managerial and professional workers.

3 The Problem

Because THS employees are hired and paid by the THS firm, they are not on the payroll of the firm actually *using* their labor. As a consequence, THS workers are not included in the employment measures generated by the BLS in its establishment-based surveys.[3]

To define this problem more formally, let $y_t = 1$ denote the event that an individual is a THS worker in period t. Hereafter the time subscript is omitted for notational convenience. The parameter of interest is the proportion of THS employees actually working in industry i, e.g., in manufacturing. That is, the probability that an individual working in industry i is a THS worker. We denote this probability by $\theta_i \equiv P(y = 1|i)$.

Our approach to estimating θ_i is quite simple. Note that this conditional probability can be written as

$$\theta_i = P(y = 1|i) = P(i|y = 1)\frac{P(y = 1)}{P(i)} \tag{1}$$

Equation (1) means that if a THS worker has a 30 percent chance of working in manufacturing or $P(i|y = 1) = 0.3$, then the number of THS workers employed in manufacturing relative to total employment–the denominator in the estimate of $P(y = 1)$−is 0.30 times $P(y = 1)$. Given that the THS industry comprised about 2 percent of the working force in 1997, and using 0.15 as the value for $P(i)$ in 1997, the proportion of THS workers relative to manufacturing employment in 1997 should be about 4 percent.

This numerical example shows that in order to estimate θ_i we need to estimate $P(y = 1), P(i)$ and $P(i|y = 1)$. The first two probabilities can be estimated by the observed proportions of THS and industry i workers from the BLS payroll surveys.

The last probability, $P(i|y = 1)$, is called the *assignment probability* because it gives the probability that a THS employee works in industry i. As mentioned in section

[3]They are, of course, included in the employment figures of the *offering* (THS) firms.

2 and in Estevão and Lach (1999), the dramatic changes in the THS industry since the late 1970s suggest that it is important to let the assignment probability vary over time. However, in trying to estimate the assignment probability on a yearly basis we face the problem that there is no direct and systematic evidence on where THS employees are actually working.[4]

Estimating $P(i|y=1)$ is the key to estimating θ_i and is, therefore, the focus of our methodology. A naive way to proceed is to assume that THS status and industry affiliation are independent. In this case, $P(i|y=1) = P(i|y=0) = P(i)$ and therefore the estimator of θ_i is the proportion of THS workers in the whole economy, $\theta_i = P(y=1)$ which does not vary across industries. Clearly, independence is too strict an assumption to make. It can, however, be relaxed a bit as follows.

Given a vector X of discrete variables characterizing the individual (education, location, occupation, gender, etc.), the assignment probability can be written as

$$P(i|y=1) = \sum_x P(i|y=1,x)P(x|y=1) \tag{2}$$

where the sum is over all possible values of X.

If we now assume that y and industry affiliation are independent *conditional* on $X = x$, (2) simplifies to[5]

$$P(i|y=1) = \sum_x P(i|x)P(x|y=1) \tag{3}$$

The assignment probability varies across industries because of differences in the distribution of characteristics across them, $P(i|x)$. Note that using data from the Current Population Surveys we can estimate the distribution of characteristics among THS (actually PSS) workers, $P(x|y=1)$, and the assignment probability of individuals having $X = x$, $P(i|x)$. Thus, $P(i|y=1)$ can be easily estimated under the conditional independence assumption.[6]

---

[4]The only direct evidence on the distribution of THS workers by industry of use appears in the Contingent Worker Supplements to the Current Population Survey in February 1995 and February 1997. See section 5.

[5]This assumption means, for example, that the probability of working in manufacturing among all female electrical engineers in Louisiana is the same irrespective of their THS status.

[6]Another possibility is to arbitrarily assign an estimate to $P(i|y=1,x)$ in (2). For example, Segal and Sullivan (1995) use $X = occupation$ and assume that all (and only) blue-collar THS employees work in manufacturing. This means assuming $P(i|y=1,x=blue\ collar) = 1$ and $P(i|y=1,x=other\ occupation) = 0$ implying $\theta_i = \frac{P(y=1,x=blue\ collar)}{P(i)}$ which is straightforward to compute.

Indeed, equation (3) reflects what is perhaps the most intuitive way of tackling the problem: estimate the distribution of some characteristic X in industry i, $P(i|X)$ and then weight it by the distribution of such characteristic X among THS workers, $P(X|y=1)$.

Even though conditional independence between industry affiliation and THS status is a weaker requirement than unconditional independence, it is still a very strong assumption to make. We will now show that in order to learn something about the time series behavior of $P(i|y=1,x)$, and therefore θ_i, we do not need to resort to drastic independence assumptions. In fact, we will show how to bound these probabilities in a non-trivial manner without making further assumptions on the relationship between industry affiliation and THS status.

4 Bounds on the Assignment Probability

We start by noting that for any value x of X, the conditional assignment probability can be written as,

$$P(i|y=1,x) = \frac{P(i,y=1|x)}{P(y=1|x)}. \tag{4}$$

We first provide bounds for the numerator of (4). Conditional on $X=x$, the probability of the joint event "the individual works in industry i" and "$y=1$" is lower than the marginal probability of each single event, i.e.,

$$P(i,y=1|x) \leq Min\left\{P(i|x), P(y=1|x)\right\}$$

implying the following upper bound for the conditional assignment probability

$$P(i|y=1,x) \leq Min\left\{\frac{P(i|x)}{P(y=1|x)}, 1\right\} \equiv \alpha_U(x) \tag{5}$$

In addition, $P(i \cup y=1|x) = P(i|x) + P(y=1|x) - P(i,y=1|x) \leq 1$ implies

$$P(i,y=1|x) \geq Max\left\{0, P(i|x) + P(y=1|x) - 1\right\}$$

and this inequality implies the following lower bound for the conditional assignment probability,

$$P(i|y=1,x) \geq Max\left\{0, \frac{P(i|x)+P(y=1|x)-1}{P(y=1|x)}\right\} \equiv \alpha_L(x) \tag{6}$$

These bounds and equation (2) prove the following proposition,

Proposition 1

$$\sum_x \alpha_L(x)P(x|y=1) \leq P(i|y=1) \leq \sum_x \alpha_U(x)P(x|y=1)$$

where $\alpha_L(x)$ and $\alpha_U(x)$ are defined in (6) and (5) and the sum is over all possible values of X.

It should be emphasized that the intervals generated by these bounds are not confidence intervals in the statistical sense. Provided the bounds are known, the interval covers the true probability $P(i|y=1)$ with probability one.[7]

A natural question to ask is whether $\alpha_L(x)$ and $\alpha_U(x)$, the bounds on the assignment probabilities, are informative. The answer depends on the choice of the conditioning vector X.

For a given value $X=x$, the lower bound on the (conditional) assignment probability is strictly positive when

$$P(i|x)+P(y=1|x) > 1 \qquad \text{(Condition L)}$$

while the upper bound is strictly less than one when

$$P(y=1|x) > P(i|x) \qquad \text{(Condition U)}$$

These conditions are more likely to be satisfied the larger $P(y=1|x)$ is. In fact, when these two conditions are satisfied the lower and upper bounds for the conditional assignment probabilities at that given value of X are

[7] Of course, when the bounds are estimated the probability will be less than 1.

$$\alpha_L(x) = \frac{P(i|x) - [1 - P(y=1|x)]}{P(y=1|x)}$$
$$\alpha_U(x) = \frac{P(i|x)}{P(y=1|x)}$$

Note that the distance between the bounds decreases with $P(y = 1|x)$: the closer $P(y = 1|x)$ gets to one, the tighter the interval containing the assignment probability. Thus, we would like to find an X that makes $P(y = 1|x)$ "large".

The intuition behind this requirement is that when $P(y = 1|x)$ is large, X is a "good" discriminant (or classifier) between THS and non-THS status in the sense that, given the individual has $X = x$, there is a high probability that s/he is a THS worker. At the limit, when $P(y = 1|x) = 1$ all individuals with $X = x$ are THS workers and the proportion of THS workers with $X = x$ assigned to industry i equals the proportion of individuals with $X = x$ assigned to industry i. That is, $P(i|y = 1, x) = P(i|x)$ when $P(y = 1|x) = 1$, and therefore the data can identify the conditional assignment probability at that particular value of X.[8]

As an illustrative example suppose that X represents location, and let New York be one such location. If it is known that all individuals living in New York are THS workers then, among all New York individuals, conditioning on the individual's THS status and location is the same as conditioning only on the individual's location. This implies $P(i|y = 1, New\ York) = P(i|New\ York)$.

To continue with this example, let us assume that $X =$ *location* takes only two values $X = 1$ for "New York" and $X = 0$ for "everywhere else". Let us further assume

[8] $P(i|x)$ is estimated by the proportion of individuals with $X = x$ working in industry i. Several data sets, such as the CPS, have information that identifies $P(i|x)$. More generally, a condition on X for (pointwise) identification of $P(i|y = 1, x)$ is that a particular value of X, say x_0, occurs only to THS workers, $P(y = 1|X = x_0) = 1$. Then,

$$\begin{aligned} P(i|y=1,x_0) &= \frac{P(i,y=1,x_0)}{P(y=1,x_0)} \\ &= \frac{P(y=1|i,x_0)P(i,x_0)}{P(y=1|x_0)P(x_0)} \\ &= \frac{P(i,x_0)}{P(x_0)} = P(i|x_0) \end{aligned}$$

because the identification condition $(P(y = 1|x_0) = 1)$ implies $P(y = 1|i, x_0) = 1$.

that no THS worker lives outside New York, implying $P(X=0|y=1)=0$ and $P(X=1|y=1)=1$. For this choice of X, Proposition 1 states

$$\alpha_L(1) \leq P(i|y=1) \leq \alpha_U(1)$$

which implies from (1)

$$Max\left\{0, \frac{P(i|X=1)+P(y=1|X=1)-1}{P(y=1|X=1)}\right\}\cdot\frac{P(y=1)}{P(i)} \leq \theta_i \leq$$

$$Min\left\{\frac{P(i|X=1)}{P(y=1|X=1)},1\right\}\cdot\frac{P(y=1)}{P(i)}$$

To get a sense of these bounds we graph them below as functions of $P(y=1|X=1)$ assuming that 41 percent of the individuals living in New York work in industry i, $P(i|X=1)=0.41$, and that $P(y=1)=P(i)$.

Note that condition L is satisfied for $P(y=1|X=1)>0.59$, while condition U is satisfied for $P(y=1|X=1)>0.41$. When $P(y=1|X=1)=1$ we have $\theta_i=0.41$.

An alternative way of giving further intuition to the bounds on the assignment probabilities is by considering the following 2×2 contingency table associated with the example above,

All individuals with $X=1$	$y=0$	$y=1$	Total
Industry i	$N_{i,y=0}$	$N_{i,y=1}$	N_i
Other industries $(\bar{i})$	$N_{\bar{i},y=0}$	$N_{\bar{i},y=1}$	$N_{\bar{i}}$
Total	$N_{y=0}$	$N_{y=1}$	$N=N_i+N_{\bar{i}}$

The population consists of N individuals with $X = 1$ and the entries in the table are frequency counts. We observe $N_{i,y=0}, N_{\bar{i},y=0}, N_{y=0}, N_{y=1}$ and N. We are looking for bounds on $P(i|y = 1) = P(i|y = 1, X = 1) = \frac{N_{i,y=1}}{N_{y=1}}$. That is, bounds on $N_{i,y=1}$.

From the table is clear that

$$
\begin{aligned}
P(i|y &= 1) = \frac{N_{i,y=1}}{N_{y=1}} \square \, Min\left\{\frac{N_i}{N_{y=1}}, 1\right\} \\
&= Min\left\{\frac{P(i|X=1)}{P(y=1|X=1)}, 1\right\} = \alpha_U(1)
\end{aligned}
$$

$$
\begin{aligned}
P(i|y &= 1) = \frac{N_{i,y=1}}{N_{y=1}} \geq Max\left\{0, \frac{N_{y=1} + N_i - N}{N_{y=1}}\right\} \\
&= Max\left\{0, \frac{P(y=1|X=1) + P(i|X=1) - 1)}{P(y=1|X=1)}\right\} = \alpha_L(1)
\end{aligned}
$$

using $N_i + N_{y=1} - N_{i,y=1} = N_i + N_{\bar{i},y=1} \square \, N_i + N_{\bar{i}} = N$ for the lower bound.

There is, of course, no guarantee that an X satisfying $P(y = 1|x) = 1$ for some x exists. In fact, it is hard to think of individual characteristics that perfectly discriminate among THS and non-THS workers. When such an X does not exist the bounds in Proposition 1 still allow us to extract some information on the assignment probability, and therefore on θ_i, without resorting to additional assumptions provided conditions L and/or U are satisfied for some values of X.

The previous discussion showed that the choice of the conditioning variables X is crucial for the quality of the bounds. What Proposition 1 is telling us is that we can choose any X (a scalar or a vector of any dimension, with as many discrete values as one wishes) and construct the bounds according to the formula in Proposition 1. In principle, one could systematically search the entire data for all combinations of conditioning variables, compute the bounds and associated intervals containing θ_i for each such vector and then take as the final interval for θ_i the intersection of all such intervals.

To conclude this section, we remark that there is a trade-off between the strength of the assumptions we are willing to make on the relationship between THS status and industry of assignment and the nature of the information on θ_i that we can extract from the data, namely, a point versus an interval estimate.

5 The Data and Estimation Issues

The Current Population Survey (CPS) is a household-based survey providing information on households' and individuals' characteristics. It assigns each individual to an industry of employment, broadly equivalent to a 3-digit SIC industry. The CPS, therefore, does not identify individuals working in the Temporary Help Supply industry (SIC 7363), but in the 3-digit industry SIC 736 that contains THS, i.e., the Personnel Supply Services (PSS) industry.[9] We use the CPS data to estimate $P(y = 1|X)$ (and $P(i|X)$). Thus, the estimates of the bounds on the assignment probability actually refer to the proportion of *PSS employees* working in manufacturing.

To be consistent with this definition, we use the level of *PSS employees* (and hours) in the economy taken from the Current Employment Survey (CES)–the establishment-based survey used in the official productivity calculations–in the calculation of $P(y = 1)$ in equation (1). The resulting estimate for the proportion of *PSS employees* among manufacturing workers should be nearly identical to the proportion of *THS employees* among manufacturing workers (θ_i). The reason is that multiplying the PSS assignment probability by the level of PSS workers gives the *flow* of PSS workers going to manufacturing. This flow, however, is about equal to the *flow* of THS workers because non-THS establishments within the PSS industry are not likely to provide workers to manufacturers since they act mostly as "matchmakers".[10] Finally, we also use the CES data to estimate $P(i)$ appearing in (1).[11]

---

[9]Figure 1 also shows the relationship between the PSS series and the Help Supply Services series (SIC 7363). The latter is a slightly broader category than purely THS firms: their difference in employment is of trivial size. Prior to the 1987 revision of the Standard Industrial Classification scheme, THS firms were classified as SIC 7362 and were part of SIC 736. SIC 736 also included employment agencies (SIC 7361) and a residual category. The 1987 revision combined the THS firms and the residual category (excluding facilities and continuing maintenance services).

[10]Non-THS establishments (SIC 7361) within the PSS industry are: chauffeur registries, maid registries, model registries, nurses' registries, ship crew registries, teachers' registries, and employment agencies. The share of non-THS establishment in total PSS employment was less than 15 percent in 1997. Non-THS employment does not contribute much to variations in PSS employment over time.

[11]A well documented discrepancy exists between the CPS and CES reported levels of PSS (SIC 736) employment, with the CPS figures below the CES numbers (Houseman and Polivka, 1998). The CES figures are taken as more accurate than the CPS ones because they are based on data provided directly by the PSS establishments. The CPS, being a household survey, may be underestimating the true number of PSS jobs because of misreporting by individuals. But correcting for this source of bias still leaves a large discrepancy between the two series. Another possible explanation for the gap between the two series is that the survey asks about *primary* jobs only while many PSS workers hold more than one job. In any case, we assume that the discrepancy between the CPS and CES figures does not bias

The top chart in figure 1 plots the personnel supply services (SIC 736) series from the CES, while the bottom chart plots the share of PSS and reported manufacturing employment. The ratio of PSS to manufacturing employment increased from about 0.01 in 1972 to 0.16 in 1997. This 16-fold increase tells much of the story behind the surge in PSS (in fact, THS) employment by manufacturers. As can be observed in equation (1), the increase in $\frac{P(y=1)}{P(i)}$ translates directly into an increase in θ_i holding the assignment probability constant.

Associated with the CPS, we also have data from the Contingent Worker Supplements to the CPS surveyed in February 1995 and in February 1997. In these supplements, respondents are asked directly if they are paid by a temporary help supply agency. Furthermore, the supplement also records their industry of assignment. Thus, these surveys constitute the only *direct* evidence on the distribution of THS workers by industry of use, $P(i|y=1)$. In 1995, 33 percent of the THS respondents worked in manufacturing but this proportion decreased to 28 percent in 1997. Redefining these proportions in terms of PSS employment, we find that in 1995 and 1997 about 27 and 22 percent of PSS employees worked in manufacturing firms, respectively. Except for these two years we do not know where PSS (THS) employees are actually working.

Before we proceed to estimate the bounds described in Proposition 1, we need to deal with four issues. First, consistent estimation of the bounds requires that sample proportions be consistent estimators of the corresponding probabilities appearing in the bounds. This is possible when individuals in each cell defined by the values of X have the same underlying probabilistic model of choice. In other words, X should capture as much of the heterogeneity as possible across individuals when making their decision to be a THS worker. This is a good reason for using a high-dimensional vector of characteristics X.[12]

The second issue is potentially more fundamental. The reported proportion of individuals in the CPS with $X = x$ working in industry i underestimates $P(i|x)$ because it counts only non-THS workers employed in industry i; it is omitting the number of

our estimates of $P(i|X)$ and $P(y=1|X)$ based on the CPS. This is certainly true if the reasons for the discrepancy are unrelated to the industry of assignment.

[12] A statistical reason, however, against using a large vector X is that the cells will have a lower number of observations than those defined by a lower dimensional X. Thus, there is a trade-off between the precision of the estimated probabilities and the bias that results from not controlling for potential heterogeneity across individuals. The balance between precision and bias in the estimation procedure is an empirical issue to be decided case by case.

THS workers in industry i. Because $P(i|x)$ enters into the calculation of the bounds we need to resolve this problem. To see this more clearly write,

$$\begin{aligned} P(i|x) &= P(i, y=0|x) + P(i, y=1|x) \\ &= P(i, y=0|x) + P(i|y=1, x)P(y=1|x) \end{aligned}$$

The data identify only $P(i, y = 0|x)$. We need to add $P(i, y = 1|x)$ but in order to correct for this "omission bias" we need to know the assignment probability $P(i|y = 1, x)$, which is precisely what we want to estimate. Not correcting for this bias underestimates both bounds on the assignment probability. To correct for this problem we used an iterative approach that takes into consideration the omission bias explicitly.[13] The final results, however, were nearly identical to the estimates assuming the absence of such bias.

The third issue arises when individuals actually employed by a THS agency report industry i as their employer, i.e. they misreport. Then the count of THS individuals underestimates the true number of THS workers and, consequently, the sample proportion of THS workers is a downward biased estimate of $P(y = 1|x)$. This source of bias can also be corrected.

Specifically, among all individuals in the CPS reporting non-THS status ($y = 0$) some are probably THS workers that did not answer the question correctly because of some confusion about who was their actual employer: the industry that uses their labor (e.g., manufacturing) or the industry that paid their wages (the THS industry). On the other hand, it is reasonable to assume that individuals reporting THS status ($y = 1$) do not make mistakes.

Let y^* represent the true THS status while y continues to denote observed status. Misreporting has implications for estimation of the bounds on $P(i|y = 1)$ because these bounds require estimates of $P(y^* = 1|x)$ and $P(x|y^* = 1)$. With misreporting, the data identify only $P(y = 1|x)$ and $P(x|y = 1)$ which differ from $P(y^* = 1|x)$ and $P(x|y^* = 1)$,

[13] At the initial stage the bounds for the assignment probabilities were estimated using the sample proportion $\frac{N_{i,y=0,x}}{N_x}$, in obvious notation, to estimate $P(i|x)$ for each x. Then a first pair of bounds, $\alpha_L^{(1)}(x)$ and $\alpha_U^{(1)}(x)$, were estimated using (6) and (5). We then used the mid-point of the estimated interval for the assignment probability–$\overline{\alpha}^{(1)}(x) = \frac{\alpha_U^{(1)}(x) - \alpha_L^{(1)}(x)}{2}$–to reestimate $P(i|x)$ using $\frac{N_{i,y=0,x}}{N_x} + \overline{\alpha}^{(1)}(x)\frac{N_{y=1,x}}{N_x}$ and derived new bounds for the assignment probabilities, $\alpha_L^{(2)}(x)$ and $\alpha_U^{(2)}(x)$. If the mid-point of the new interval differed significantly from its previous value we reestimated $P(i|x)$ and obtained a third round of estimated assignment probabilities. This process was iterated until (pointwise) convergence of $\overline{\alpha}^{(n)}(x)$, where n is the iteration stage.

$$
\begin{aligned}
P(y^* = 1|x) &= P(y^* = 1, y = 1|x) + P(y^* = 1, y = 0|x) \\
&= P(y^* = 1|y = 1, x)P(y = 1|x) + P(y^* = 1|y = 0, x)P(y = 0|x) \\
&= \underbrace{P(y = 1|x)}_{observed} + \underbrace{P(y^* = 1|y = 0, x)}_{misreporting\ rate} P(y = 0|x)
\end{aligned} \tag{7}
$$

assuming that reporting THS individuals do not make mistakes.

Conditional on $X = x$, the true proportion of THS workers equals the observed proportion of THS workers plus a percentage–the misreporting rate–of the observed proportion of non-THS. Clearly,

$$P(y^* = 1) = \sum_x P(y^* = 1|x)P(x) \tag{8}$$

and,

$$P(x|y^* = 1) = P(y^* = 1|x)\frac{P(x)}{P(y^* = 1)} \tag{9}$$

It follows from (7), (8) and (9) that given an estimate of the misreporting rate, i.e., the probability that a reporting non-THS individual is in fact a THS worker, we can adjust all observed proportions to form consistent estimates of the true probabilities.

Fortunately, the misreporting rate can be estimated from the Contingent Worker Supplements to the CPS of February 1995 and February 1997. The supplement asks each worker in the main CPS whether he was hired by a temporary help supply firm or not. By comparing the worker's answers in the supplement to the answers in the main CPS, we can estimate $P(y^* = 1|y = 0, x)$ for these years.

Because we do not have information on other years we will assume that the misreporting rate in every year other than 1995 and 1997 is equal to the average of the misreporting rates in 1995 and in 1997. Note, however, that the probability of misreporting is conditional on $X = x$. To the extent that the mix of characteristics varies across years, we may get significant variation in the degree of misreporting in each year.

Having estimated $P(y^* = 1|x)$ for every value of X according to (7), we average them using CPS population weights to estimate $P(y^* = 1)$ using (8). Given these two estimates, we use (9) to estimate $P(x|y^* = 1)$. Finally, we use the estimates of $P(y^* = 1|x)$ and $P(x|y^* = 1)$ to estimate the bounds in (6) and (5).

The last point is related to the "omission bias" mentioned above. Observe that the number of workers in industry $i-$the denominator in the definition of θ_i in (1)–should be the *true* number of workers in industry i, that is, the reported number plus the THS employees working in the industry. Thus, given an estimate of the assignment probability $\widehat{P}(i|y=1)$ (a point estimate if available or the bounds in Proposition 1), the estimated probability of finding a THS worker in industry i is

$$\widehat{\theta}_i = \frac{\overbrace{\widehat{P}(i|y=1)N_{y=1}}^{N_{y=1,i}}}{N_{y=0,i} + \underbrace{\widehat{P}(i|y=1)N_{y=1}}_{N_{y=1,i}}} \tag{10}$$

where $N_{y=1}$ is the observed number of THS workers and $N_{y=0,i}$ is the reported number of industry i workers. $N_{y=1,i}-$the numerator in (10)–is the number of THS workers in industry i while the denominator is the (true) number of workers in industry i including THS workers, $N_i = N_{y=0,i} + N_{y=1,i}$.

We estimated the bounds for the proportion of THS workers in manufacturing using data from the March CPS tapes for the 1972-1997 period.[14] We also estimated the proportion of THS hours in manufacturing following the same methodology.[15]

After some experimentation with different conditioning variables we decided on a conditioning vector X that includes occupation (2-digit), state of residence, educational achievement, gender, age and a dummy variable for whether the individual works part-time or not.[16]

For our choice of X, the estimates of the conditional probabilities $P(i|X)$ and

[14]The universe is defined as employed workers who, at the time of the survey, are not self-employed and do not work in farms, fishing or forestry. Each individual observation is weighted by its sample population weight.

[15]For hours, we weighted each individual observation used in computing the bounds on the assignment probabilities by the average hours worked by the individual during the week before the survey was taken. Because there is no direct information on average weekly hours of PSS workers (SIC 736), $N_{y=1}$ in (10) refers to the number of PSS hours obtained by multiplying PSS employment by average weekly hours for help supply workers (SIC 7363).

[16]Excluding any of these variables makes the bounds slightly wider and including extra variables affects the bounds just marginally. The variables that contributed most to the narrowing of the distance between both bounds were, in order of importance, occupation, educational achievement and state of residence. We avoided using more finely defined variables, such as occupation at the 3-digit level of aggregation, because this reduces considerably the number of observations in some cells. We also experimented with coarser breakdowns of the occupation and education variables–12 and 5 categories in 1985, respectively–while keeping the remaining variables untouched. As expected, the bounds tend to be tighter once finer breakdowns of each variable are allowed.

$P(y = 1|X)$ were generally based on cells with a small number of observations. For this reason, we computed confidence intervals around the bounds (derived in the Appendix). In order to guarantee the feasibility of the computation of the confidence intervals, we did not use the cells–realizations of the vector X in which there was only one observation (individual). This left us with about 8,000-13,000 cells, depending on the year.

Note also that the bounds on $P(i|y = 1)$, and therefore on θ_i, are estimated independently for every year. Thus, the underlying function relating individual characteristics to the probabilities of assignment is allowed to change over time, as suggested by the description of the THS industry in section 2.

6 Estimates of Manufacturers' Use of THS Workers

Table 1 shows the estimated bounds, and the mid-point of the interval for θ_i. Estimates for the proportion of THS hours in manufacturing begin in 1982 because the BLS payroll survey does not provide information on the average hours of work in this industry for the earlier period. A number of interesting points are worth emphasizing.

First, the bounds are quite informative. Second, both bounds exhibit an upward trend over time as would have been expected from the anecdotal evidence on the increasing use of temporary help supply arrangements discussed in the introductory section. While in the first years of our sample the lower bound for the probability of finding a THS worker in manufacturing was negligibly different from zero, after 1991 the lower bound is consistently above 3/4 percent. The upper bound also presents a clear upward trend, going from about 2/3 percent in the first half of the 1970s to above 5 percent in the late 1990s. In fact, the average lower bound for θ_i in the 1995-97 period is larger than the average upper bound at the end of the 1970s, strongly suggesting a regime shift in manufacturers' hiring patterns in the 1980s and 1990s.

This trend can be seen more clearly in figure 2 where the bounds and the midpoint of the interval between the lower and upper bounds are plotted. Asymptotic 95 percent confidence intervals are quite tight and cannot be even noticed in the lower panel of figure 2.

Third, the bounds exhibit a cyclical pattern that is consistent with the idea that manufacturers use this form of employment as an adjustment margin for sudden economic shocks. For instance, as a response to the economic slowdown, the use of THS workers decreased significantly in 1991. Subsequently, from 1992 onwards manufacturers stepped up hiring of THS workers while leaving payroll employment below its 1991 level until 1995 (see discussion of table 2 in section 8.)

Figure 3 presents the bounds on the assignment probability ($P(i|y=1)$) as per Proposition 1. Notice that the bounds here refer to the proportion of *PSS employees* going to manufacturing. Unlike the bounds for θ_i, these bounds exhibit a less discernible trend. The mid-point between the bounds decreases slightly from about 29 percent at the beginning of the sample to 23 percent at the end of it–about the same value obtained from the Contingent Worker Survey for 1997. The bounds are also tighter in later years, suggesting that the vector of worker characteristics, X, is a better indicator of the industry of assignment in the latter portion of our sample. The difference between the upper and the lower bounds at the beginning of the sample is about 50 percentage points whereas, in the last two years of the sample, this difference declined to 30 percentage points.

The relative stability of the assignment probability over time confirms that the increase in θ_i is the result of the 16-fold increase in $\frac{P(y=1)}{P(i)}$ (see figure 1). That is, about 25-30 percent of PSS workers are absorbed by manufacturers every year. Over time, however, there are substantially more THS manufacturing workers and less non-THS manufacturing workers. As a result the share of THS workers in manufacturing increased dramatically.

The reported bounds give us already a substantial amount of information on θ_i. Nevertheless, it is sometimes useful to have point estimates of the parameters of interest. We will, in fact, use point estimates to address measurement issues in section 8. We can compute three point estimates based on different types of "identifying assumptions" about the assignment probability.

First, in the absence of additional information, the mid-point between the bounds is a reasonable choice for a point estimate.[17] As mentioned in section 3, a second point estimate can be derived from the identifying assumption that $P(y=1|x,i) = P(y=1|x)$, i.e., the events "being a THS worker" and "working in manufacturing" are independent

[17]We can offer two justifications for using the midpoint of the interval as a point estimate of θ_i. First, the midpoint is the value of θ_i that minimizes the maximal error. More precisely, if we choose any point z in the interval $[\theta_{iL}, \theta_{iU}]$, the maximal error is $Max\,\{z-\theta_{iL}, \theta_{iU}-z\}$. The value of z that minimizes this maximal error is the midpoint.

In addition, suppose $\widehat{\theta}_i$ is an unbiased estimator of θ_i, but θ_i itself is also a random variable. The expected value of $\widehat{\theta}_i$ conditional on a realized value of θ_i is $E(\widehat{\theta}_i|\theta_i) = \theta_i$, and the unconditional expectation is $E(\widehat{\theta}_i) = \int_{\theta_{iL}}^{\theta_{iU}} \theta_i f(\theta_i) d\theta_i$,where $f(\theta_i)$ is the density function of θ_i. It is straightforward to prove that if $f(\theta_i)$ is symmetric in the interval $[\theta_{iL}, \theta_{iU}]$ then $E(\widehat{\theta}_i) = \theta_{iL} + \frac{\theta_{iU}-\theta_{iL}}{2}$.

once we control for a vector of worker characteristics (see equation (3)). We call this estimator the "conditional independence" estimator.

A third point estimate is obtained by making the identifying assumption that a particular realization of X is sufficient to perfectly discriminate among the industries of assignment. One such assumption made by Segal and Sullivan (1995) is that all (and only) blue-collar PSS workers are employed by manufacturers. Setting $X = occupation$ and assuming $P(i|y = 1, x = blue\ collar) = 1$ and $P(i|y = 1, x = other\ occupation) = 0$ implies $\theta_i = \frac{P(y=1,x=blue\ collar)}{P(i)}$ which is straightforward to compute (see equations (1) and (2)). We call this estimator the "blue collar" estimator.

All three point estimates are plotted in figure 4. All estimates suggest an upward trend in the hiring of THS workers by manufacturers. There is even the suggestion of a trend break at the beginning of the 1990s. As mentioned before, their movements over time are consistent with the hypothesis that manufacturers have been using THS workers as an adjustment margin to economic shocks.

Note that the "conditional independence" estimator mimics the behavior of the mid-point estimator pretty well, except that it is shifted downward. The "blue-collar" estimator suggests that manufacturers were not using nearly as much THS workers in the 1970s and in the 1980s as implied by the previous point estimates. This prompts a more dramatic increase in manufacturers' use of THS workers in the 1990s. It should be pointed out, however, that some of the annual estimates in the 1980s fall below the theoretical lower bound casting serious doubts on the validity of this particular identifying assumption. The rejection of the "blue-collar" estimator shows that the bounds are useful to assess the validity of particular point estimates.

7 Sensitivity Analysis

The bounds on the proportion of THS employees working in manufacturing are derived from a theoretical construct. We just saw that their pattern over time conforms with prior expectations, but is the level of the implied proportions sensible? Can they be off the mark?

To answer these questions, we bring additional information on the THS industry which we use to estimate θ_i for selected years. If these estimates fall within the bounds, we will take this as supporting evidence for our methodology.

Our first set of estimates is derived from the estimated assignment probabilities

$P(i|y=1)$ taken directly from the Contingent Worker Supplement (CWS) to the CPS of February 1995 and February 1997. Given these $\widehat{P}(i|y=1)$s, we estimated θ_i for manufacturing using equation (10). The probability of finding a THS worker in manufacturing in 1995 and in 1997 was about 3.4 percent. The probability of finding an hour of THS work in the manufacturing sector was about 2.8 percent in 1995 and a bit more in 1997.

A second set of estimates may be obtained by using information from the input-output tables provided by the Bureau of Economic Analysis every five years. When wages of PSS workers and other fees are independent of their industry of assignment, the proportion of the PSS industry's output that goes to manufacturing–the input-output coefficient–may be used as estimates for the proportion of PSS hours used by the manufacturing sector. Input-output tables with the relevant information on the PSS industry are available for 1977, 1982, 1987 and 1992.[18]

Figure 5 plots the CWS and the input-output estimates of θ_i: They are all contained within the estimated theoretical bounds even though their underlying data come from independent sources. The estimation problems associated with each set of estimates notwithstanding, this finding is certainly comforting. In particular, the CWS estimates are nearly identical to the mid-point between the bounds for 1995 and 1997. This result certainly corroborates our expectation that the mid-points provide reasonable information on the parameter of interest when no other information is available.

We can also use the CWS data for February 1995 and February 1997 on THS workers' characteristics and information on where they are actually working to shed some preliminary light on the determinants of the assignment probability $P(i|y=1)$. A logit regression of the binary variable $y=1$ on individual characteristics suggests that well-educated, blue-collar THS workers have a higher probability of working in manufacturing. The latter also increases for THS workers between 35 and 50 years of age, for those living in the Midwest and for those living outside metropolitan areas. Finally, the probability of being a THS worker in manufacturing declines if the individual works part-time.[19]

[18]The output of the PSS sector can be written as $Y = w_m H_m N_m + w_r H_r N_r$, where the subscript $i = m, r$ indicates the industry of assignment (m = manufacturing and r = non-manufacturing), w_i = hourly wage plus hourly overhead fees, H_i = average hours of work and N_i = number of workers assigned to industry i. If $w_m = w_r$, then the proportion of PSS output going to manufacturing (the input-output coefficient) is the share of total hours of work going to manufacturing. However, we do not have information on the evolution of the gap between $w_m$and w_r and the estimates of θ based on the input-output tables should be viewed with caution.

[19]Detailed results appear in the working paper version of the paper and are not presented here to

Actually, the estimated coefficients in the logit regression can be used to forecast the probability that a THS employee works in manufacturing using data on PSS workers' characteristics from the March CPS tapes in each year between 1972 and 1997. The estimate of $P(i|y=1)$ is the weighted average of the individual predicted probabilities using CPS population weights. $\widehat{\theta}_i$ is calculated by plugging the estimates of $P(i|y=1)$ into equation (10) in every year.[20] The predicted CWS-CPS proportions of THS employees and hours in manufacturing using the March CPS files all fall within the estimated bounds lending further support to our methodology.

8 Implications for Employment and Labor Productivity in Manufacturing

As suggested in section 6, in the absence of additional information, we can use the mid-point estimate of the interval generated by the bounds on θ_i. Then, combining the mid-point estimates for the period 1972-1994 and for 1996 with direct information on the industry assignment of THS workers in 1995 and 1997, we get a complete time series for θ_i. With such an estimated time series for θ_i in hand we can provide quantitative answers to several interesting questions.

8.1 How many people actually work in manufacturing after accounting for THS?

Let N_i and $N_{i,y=0}$ be the true and the reported number of workers in manufacturing respectively. Then,

$$N_i = N_{i,y=0} + \theta_i N_i$$

conserve space. The regressors are occupation, educational level, race, gender, age, part-time status, region of residence and a binary variable for living in a metropolitan area. To increase the sample size, we pooled the 1995 and 1997 surveys and included a dummy variable for 1997.

[20] We should note that the logit forecasting exercise is based on very strong assumptions. First, the discrete choice model needs to be correctly specified both in terms of the variables affecting the assignment probability and in terms of functional form. Second, we need to make the crucial assumption that the parameters of the model remain unchanged over time. This is a problematic assumption in view of the dramatic changes experienced by the THS industry at the beginning of this decade. Third, the discrete choice approach is limited to high industry aggregation levels because it will be extremely difficult to pursue at more disaggregated levels due to the lack of enough observations per industry. Finally, the model is fitted using the universe of *THS* employees but the assignment probabilities are forecasted using data on *PSS* employees.

$$\implies N_i = \frac{N_{i,y=0}}{1-\theta_i}$$

A similar formula is used for hours. Figure 6 plots the reported series of manufacturing employment ($N_{i,y=0}$) and hours from the CES and the respective series adjusted for THS employment and hours. Columns 1 and 3 of table 2 report the data underlying figure 6, while the second column displays the time series of THS workers in manufacturing ($\theta_i N_i$) using the estimated time series for θ_i.

Between 1991 and 1997 the number of THS workers in manufacturing expanded at an annual rate of 11 percent. As a consequence, manufacturing employment adjusted for THS workers grew at about 0.5 percent at an annual rate whereas reported employment grew only at half this rate. Put it differently, between 1991 and 1997, manufacturing payroll employment increased by about 271,000 persons, but the inclusion of THS workers elevates this figure to 580,590. Hours worked exhibit a similar pattern.

In addition, accounting for THS workers, manufacturing employment returned by 1997 to within 450,000 workers of its 1989 level (a peak year) in contrast to CES figures that place manufacturing employment at 715,000 workers short of the 1989 level in 1997. Moreover, the decline in manufacturing hours between the local peak in 1989 and 1997– about 1.25 percent–virtually disappears once THS workers are taken into account.

These findings are consistent with the view that the THS industry has been facilitating rapid changes in the level of employment in firms that otherwise would be more reluctant to change their permanent labor force in the face of what may be temporary changes in demand conditions. Furthermore, the level of THS employment in manufacturing at the end of our sample is significantly high by historical standards, suggesting that, in the event of a downturn, much of the adjustment in labor input can be accomplished with small effects on the reported manufacturing payroll employment series.

8.2 How variable, and cyclical, is THS employment in manufacturing?

The year-to-year variation in THS employment and hours is of an order of magnitude larger than for non-THS. The coefficient of variation of non-THS manufacturing employment is about 0.04; the coefficient of variation for manufacturing THS employment is 0.7.

The estimates of THS workers suggest a noticeable degree of cyclicality in manufacturers' use of THS hours and employment. The wedges between each adjusted and the unadjusted series are smaller during the downturn of the beginning of the 1990s and widen during the subsequent expansion. These estimates are also consistent with the anecdotal evidence pointing to a large increase in manufacturing THS employment at the end of the 1980s.

8.3 How does the use of THS workers affect the measurement of labor productivity?

The differences in the levels and growth rates between the reported data and the data adjusted for the use of THS workers imply that reported labor productivity is overstated in the manufacturing sector and understated in the service sector.

Columns 3 and 8 in Table 3 show the official (BLS) statistics for labor productivity growth. The bottom part of the table illustrates the puzzle that has been voiced by policy makers, academics and the popular press. The trend growth in manufacturing labor productivity is significantly larger than the trend productivity growth in the whole non-farm business sector. For instance, during the period 1982 to 1997, manufacturing labor productivity grew at an annual rate of 3.2 percent while labor productivity in the non-farm business sector grew at an annual rate of 1.8 percent. Furthermore, while aggregate productivity growth has not changed much since the 1980s, manufacturing labor productivity has accelerated substantially. How much of this wedge can be accounted for by the mismeasurement of the labor input?

Column 4 reports all persons hours adjusted for hired THS hours in manufacturing. Comparing columns 3 and 5 gauges the effect on measured productivity growth of adjusting for THS hours in manufacturing. As expected, the effect is strongest during the 1990s when the use of THS workers picked up. On this account, between 1991 and 1997 the average growth rate of labor productivity in the manufacturing sector was biased upward by 0.27 percentage point per year. In other words, adjusting for the increase in the use of THS workers lowers the measured growth rate of productivity during this period from about 3.9 percent per year to 3.6 percent per year.

The small differences in the growth rates accumulate over time. Assuming that θ_i is zero before 1982 and thereafter given by the mid-point estimates and the direct survey evidence for 1995 and 1997 implies that the productivity level in 1997 was about 2.5 percent lower than the reported one. Nevertheless, THS-adjusted productivity still accelerates in the 1990s and the full explanation for the divergence between manufacturing and non-farm business productivity growth lies elsewhere.[21]

9 Conclusions

The goal of this paper was to estimate the extent to which the manufacturing sector is outsourcing labor from the service sector via the hiring of THS workers.

We develop a new methodology to estimate the proportion of THS employees working in other industries. Our approach uses minimal assumptions and is non-parametric. It consists of establishing bounds for the probability that a THS worker does indeed work in a manufacturing industry. We develop conditions under which these bounds are tight and estimate them using readily available data from the March tapes of the CPS. In the same spirit of Manski's work (1995), this methodological framework may be used more generally to put bounds on conditional probabilities that are not directly observed by the researcher. These bounds are not just useful by themselves but also provide a powerful tool to assess the reliability of point estimates for the parameter under study.

In our particular application, the estimated bounds are informative and confirm that manufacturing firms have increased the use of temporary help supply workers during the 1990s. Accounting for THS workers, manufacturing employment returned by 1997 to within 450,000 workers of its 1989 level (a peak year) in contrast to BLS payroll figures that place manufacturing employment at 715,000 workers short of the 1989 level in 1997. Moreover, the decline in manufacturing hours between the local peak in 1989 and 1997–about 1.25 percent–disappears once THS workers are taken into account.

The time series behavior of the estimated series of THS employees (hours) in manufacturing is consistent with the hypothesis that manufacturers have been using this type of work as a buffer to economic shocks. However, accounting for THS employment explains only a small part of the divergence between labor productivity growth in manufacturing and elsewhere in the 1990s.

[21] Most economists would single out the mismeasurement of output outside the manufacturing sector as the most likely explanation for this productivity gap. At the core of this argument lies the suspicion that the price deflators used by the Commerce Department to get real output data from nominal output series overstates the actual inflation in the service sector and thus understates real output growth outside manufacturing. Others, e.g., David (1990), suggest that the gains from the new information technologies are yet to come: while manufacturers have appropriated some of these gains, its diffusion to the rest of the economy is slow.

References

[1] Abraham, Katherine G. and Susan K. Taylor, "Firms' Use of Outside Contractors: Theory and Evidence", *Journal of Labor Economics* 1996, 14,3, pp.394-424.

[2] Autor, David, "Why Do Temporary Help Firms Provide Free General Skills Training?", mimeo, JFK School of Government, Harvard University, October 1998.

[3] David, Paul, "The Dynamo and the Computer: A Historical Perspective on the Modern Productivity Paradox", *American Economic Review*, Papers and Proceedings, May 1990, 80(2):355-361.

[4] Estevão, Marcello and Saul Lach, "The Evolution of the Demand for Temporary Help Supply Employment in the United States", NBER WP 7427, December 1999.

[5] Golden, Lonnie "The expansion of Temporary Help Employment in the US, 1982-1992: A Test of Alternative Economic Explanations", *Applied Economics*, 28, forthcoming.

[6] Houseman, Susan and Anne Polivka, "The Implications of Flexible Staffing Arrangements for Job Security", mimeo, April 1988.

[7] International Herald Tribune, September 2, 1997.

[8] Manski, Charles, *Identification Problems in the Social Sciences*, Cambridge and London: Harvard University Press, 1995.

[9] National Association of Temporary and Staffing Services, various press releases.

[10] Polivka, Anne E., "Are Temporary Help Agency Workers Substitutes for Direct Hire Temps? Searching for an Alternative Explanation of Growth in the Temporary Help Industry", mimeograph, BLS, 1996.

[11] Segal, Lewis M. and Daniel G. Sullivan, "The Temporary Labor Force", *Economic Perspectives*, 19, 1995, 2-19.

[12] Segal, Lewis M. and Daniel G. Sullivan, "The Growth of Temporary Services Work", *Journal of Economic Perspectives*, Spring 1997, 11, 117-136.

A Appendix

In this appendix we show how to construct confidence intervals for the bounds on θ_i. First, because the number of observations in the CES is very large we treat the sample proportions of individuals working in manufacturing and in the PSS industry as the true parameters $P(i, y = 0)$ and $P(y = 1)$. The number of observations of individuals working in the PSS sector in the CPS is also sufficiently high to treat the sample proportion for PSS individuals with characteristics x as the true probability $P(x|y = 1)$.[22]

The cells defined by the different realizations of the vector X, however, do not have that many observations. Even though we excluded from computation of the bounds those cells that have a single observation, we still have many cells with 2, 3, etc. individuals. Thus, the estimates of $P(y = 1|x)$ and $P(i|x)$ may not be very precise. In this appendix we show how to build confidence intervals for $P(y = 1|x)$ and $P(i|x)$ and how to use them to construct asymptotic 95% confidence intervals for the bounds on θ_i.

Note that y_j–the binary variable indicating whether individual j having characteristic x is a THS worker or not–is a Bernoulli variable with probability $P(y = 1|x) = \pi(x)$. Similarly, i_j is the Bernoulli variable indicating whether individual j works in industry i with probability $\lambda(x) = P(i|x)$. An easy way to construct a confidence interval for these probabilities is to rely on the asymptotic distribution of sample proportions. It is straightforward to show that 95% confidence intervals for $\pi(x)$ and $\lambda(x)$ are given by

$$\left[\overbrace{\widehat{\pi}(x) - 1.96\sqrt{\frac{\widehat{\pi}(x)\left(1-\widehat{\pi}(x)\right)}{n(x)}}}^{\widehat{\pi}_L(x)}, \overbrace{\widehat{\pi}(x) + 1.96\sqrt{\frac{\widehat{\pi}(x)\left(1-\widehat{\pi}(x)\right)}{n(x)}}}^{\widehat{\pi}_U(x)}\right]$$

$$\left[\underbrace{\widehat{\lambda}(x) - 1.96\sqrt{\frac{\widehat{\lambda}(x)\left(1-\widehat{\lambda}(x)\right)}{n(x)}}}_{\widehat{\lambda}_L(x)}, \underbrace{\widehat{\lambda}(x) + 1.96\sqrt{\frac{\widehat{\lambda}(x)\left(1-\widehat{\lambda}(x)\right)}{n(x)}}}_{\widehat{\lambda}_U(x)}\right]$$

where $n(x)$ is the number of individuals having $X = x$ and $\widehat{\pi}(x)$ and $\widehat{\lambda}(x)$ are the sample proportion of those individuals who are THS workers and of those individuals working in industry i, respectively.

[22] For instance, the March CPS tapes have between 275 and 400 observations for the PSS industry in every year after 1985.

Recall that the upper bound on the assignment probability is

$$\alpha_U = \sum_x Min\left\{\frac{P(i|x)}{P(y=1|x)}, 1\right\} P(x|y=1)$$

An asymptotic 95 percent confidence interval for α_U is then

$$CI(\alpha_U) = \left[\sum_x Min\left\{\frac{\widehat{\lambda}_L(x)}{\widehat{\pi}_U(x)}, 1\right\} \widehat{P}(x|y=1), \sum_x Min\left\{\frac{\widehat{\lambda}_U(x)}{\widehat{\pi}_L(x)}, 1\right\} \widehat{P}(x|y=1)\right]$$

and, analogously, an asymptotic 95 percent confidence interval for

$$\alpha_L = \sum_x Max\left\{0, \frac{P(i|x) + P(y=1|x) - 1}{P(y=1|x)}\right\} P(x|y=1)$$

is,

$$CI(\alpha_L) = \left[\begin{array}{c} \sum_x Max\left\{0, \frac{\widehat{\lambda}_L(x) + \widehat{\pi}_L(x) - 1}{\widehat{\pi}_U(x)}\right\} \widehat{P}(x|y=1)\ , \\ \sum_x Max\left\{0, \frac{\widehat{\lambda}_U(x) + \widehat{\pi}_U(x) - 1}{\widehat{\pi}_L(x)}\right\} \widehat{P}(x|y=1) \end{array}\right]$$

It follows that an asymptotic 95 percent confidence interval for the upper bound on θ_i is, therefore, $\frac{\widehat{P}(y=1)}{\widehat{P}(i,y=0)+\widehat{\alpha}_L \widehat{P}(y=1)} \times CI(\alpha_U)$, while the corresponding confidence interval for the lower bound on θ_i is $\frac{\widehat{P}(y=1)}{\widehat{P}(i,y=0)+\widehat{\alpha}_U \widehat{P}(y=1)} \times CI(\alpha_L)$.