The long-run cost of job loss as measured by consumption changes

Browning, Martin; Crossley, Thomas F.

Accepted Manuscript

The long-run cost of job loss as measured by consumption changes

Martin Browning, Thomas F. Crossley

PII: S0304-4076(08)00056-0
DOI: 10.1016/j.jeconom.2008.05.005
Reference: ECONOM 3036

To appear in: *Journal of Econometrics*

Please cite this article as: Browning, M., Crossley, T.F., The long-run cost of job loss as measured by consumption changes. *Journal of Econometrics* (2008), doi:10.1016/j.jeconom.2008.05.005

The Long-Run Cost of Job Loss as Measured by Consumption Changes

Martin Browning

University of Oxford and IFS

Thomas F. Crossley[1]

University of Cambridge, McMaster University and IFS

January 2008

[1]Corresponding Author. Faculty of Economics, University of Cambridge, Cambridge CB39DD, England. Phone: +44-1223-335225. Fax: +44-1223-335475. Thomas.Crossley@econ.cam.ac.uk.

Abstract

We consider the measurement of the cost of job displacement. With a Canadian panel survey we compare the consumption growth of households that experienced a permanent layoff to a control group of households that experienced a temporary layoff with known recall date. Because the firms employing the latter group are providing insurance, these workers approximate a benchmark of full insurance against job loss shocks. We estimate that permanent layoffs experience an average consumption loss of between 4 and 10 percent. Older workers and workers with high job tenure have losses closer to the top of this range.

JEL Classifications: D91, J63, J65

Keywords: Job Displacement, Consumption

1 Introduction

For many workers the loss of a job because of plant closure or a permanent layoff may involve a considerable loss of lifetime welfare. These workers bear a disproportionate share of the costs of reallocation in a dynamic economy. Given the potential for large losses, there a number of alternative policies for governments to follow. One is the employment protection route in which governments make it as difficult as possible for firms to lay off workers. This has a potentially harmful impact on hiring and does not address the losses to workers when firms do in fact go bankrupt. A second option is to provide generous unemployment benefits for a long time in order to allow workers to search for the best fit in a new job. Once again this has deleterious side effects and still does not address the issue that even with such a Unemployment Insurance scheme some workers will still experience a large negative permanent shock. A third (hypothetical) policy option is to provide full insurance against such losses. In this paper we attempt to quantify the gains from such insurance.

Job displacement studies have typically been concerned with the effect of displacement on short run earnings and wages and the duration of joblessness. The attempt to quantify the long run welfare loss due displacement against a full insurance benchmark faces at least two major problems. The first of these is the difficulty in measuring changes in lifetime welfare. The second problem is that given a sample of displaced workers we do not have a natural control group who faced full insurance and hence experienced no gain or loss consequent on their being displaced or not.

As regards the first problem, even if we have long panels and examine earnings, the mapping from wage or earnings paths to lifetime welfare is not a simple one. In all but the simplest frictionless labour market models wages depend on household preferences (discount factors, risk aversion and prudence) and the possibilities open to households for intertemporal smoothing. Thus changes in wage or earnings possibilities cannot be simply mapped into changes in welfare without an explicit theory model. In addition, if there are other potential or actual earners in the household then even a large loss of earnings by one partner may not lead to a sharp fall in lifetime welfare. Thus the presence of other potential earners in the household provides some natural (self) insurance even without outside options. These difficulties are compounded by the fact that we do not usually have a long panel so that we have to extrapolate from short run changes using standard earnings processes. The latter may not be reliable for those who have recently been displaced.

In the paper we propose dealing with this problem by using changes in consumption to trace out long run impacts. Just as with wages or earnings, this requires a formal model within which we can measure and interpret the lifetime loss from observations on consumption before and after job displacement. We develop such a framework in Section 2 using a conventional life-cycle model with forward looking agents. Our framework takes account of short run adjustment to the displacement, the possible presence of other earners, other idiosyncratic shocks, macro shocks and changes in demographics. Within this framework we can define a 'treatment' (the negative economic shock associated with an imperfectly insured job loss) and an appropriate counter factual (perfect insurance against such shocks). We can also discuss the choice of estimator. As we shall see, a convenient estimator within our theoretical framework uses matching techniques.

The second major problem mentioned above is not having a natural control group. In the displacement literature the usual comparison is between displaced workers and workers who retain their job. While this may be a useful comparison in some contexts, it is not appropriate for our purposes since workers who continue in their jobs and who had a positive prior probability of losing the job actually experience a welfare *gain.* Uncertainty has resolved in their favour. In this view, the comparison of the displaced to those who retain their job overestimates the loss (relative to full insurance) experienced by the displaced. Instead we compare changes in lifetime outcomes between the displaced and those who were temporarily laid off and expect to be recalled at a specific date. We maintain that this group is closer to the desired counter-factual since they are in firms that use temporary layoffs and hence provide considerable insurance (albeit, possibly less than full) to their workers.

Naturally, the use of temporarily laid off workers to estimate the counterfactual for permanent laid off workers requires a conditional independence assumption: conditional on observables, expected consumption growth under full insurance should be the same for the two groups. The plausibility of this assumption is helped by the richness of the set of observables our data allow us to condition on.

This paper is a contribution to three literatures. The first is the extensive literature on the effects of job displacement; see and Fallick (1996), Kletzer (1998) and Kuhn (2002) for surveys. Second, this paper is related to tests of full insurance and consumption growth around idiosyncratic shocks such as job loss, illness or disability; see, for example, Cochrane (1991) and Stephens (2001). Finally, what we present here is a complement to work on the short run costs of job loss and the impact of Unemployment Insurance benefits (see Gruber (1997) and Browning and Crossley (2001) and (2003)).

Our main finding is that permanently displaced workers suffer an average consumption loss of between 4 and 10 percent. Older workers and workers with high job tenure have losses closer to the top of this range. As this estimate is relative to our best approximation of a full insurance benchmark, it provides an upper bound on the value of new policy initiatives designed to mitigate the costs of job loss.

The next section develops our theoretical framework, which in turn suggests a natural estimation strategy. Section 3 describes the data. Our results are presented in Section 4. Section 5 discusses the policy implications of our results and Section 6 concludes.

2 Theoretical Framework

In a conventional life-cycle model (which assumes a forward-looking, optimizing household), the marginal utility of expenditure *mue* λ_t evolves according to:

$$\lambda_{t+1} = \lambda_t + \varepsilon_{t+1}, E_t\left(\varepsilon_{t+1}\right) = 0. \tag{1}$$

We now develop a framework that will allow us to quantify the effect of a job loss in terms of the *mue*. Risk averse households desire to hold the *mue* constant across possible states of the world. In the ideal situation, in which society provided full insurance against idiosyncratic risk, a household's *mue* responds only to aggregate shocks.

Let d be an indicator variable that takes value 1 if the agent keeps her job from period t to $t+1$ and 0 if she is displaced. $E_t()$ denotes the expectation for a given agent, conditional on the information available to the agent at time t. We assume an additive structure for the job retention/loss shock and other shocks:

$$\varepsilon_{t+1} = (1-d)\,\Gamma_t^0 + d\Gamma_t^1 + \eta_{t+1} \tag{2}$$

where Γ_t^d is the shock consequent on the realization d and η_{t+1} is the effect of other shocks. The additivity here between job retention/loss shocks and other shocks is for convenience. Nevertheless, without some structure it is much hard to make sense of the question "what are the costs of job loss?"[1] Note that we do not assume that job loss and other shocks are independently distributed.

We shall always assume:

$$\Gamma_t^0 > 0 > \Gamma_t^1 \tag{3}$$

so that a job loss is equivalent to a wealth loss (which raises the *mue*, all other things being equal). Critically for the development below, it follows that retaining a job is a positive shock for agents who faced a positive probability of job loss (the *mue* falls, indicating that the agent is better off). If the agent had full insurance against job loss shocks then we would have:

$$\Gamma_t^0 = \Gamma_t^1 = 0 \tag{4}$$

since the realization of d does not make any difference to the agent. In the displaced worker literature attention has focussed on differences between remaining in the job and being displaced; in the current context this is given by $(\Gamma_t^1 - \Gamma_t^0)$. For our policy driven analysis the appropriate object of interest is Γ_t^0, since the full insurance benchmark is zero.

Let π_t be the probability at time t of $d = 1$. We have:

$$\begin{aligned} E_t\left(\varepsilon_{t+1}\right) &= \pi_t\left(\Gamma_t^1 + E_t\left(\eta_{t+1} \mid d = 1\right)\right) + \\ &\quad (1 - \pi_t)\left(\Gamma_t^0 + E_t\left(\eta_{t+1} \mid d = 0\right)\right). \end{aligned} \tag{5}$$

[1] Absent additivity, answering this question would requires some arbitrary decomposition of the effects of the shocks.

Combining this with (1) we have:

$$\pi_t \left(\Gamma_t^1 + E_t \left(\eta_{t+1} \mid d = 1\right)\right) = -\left(1 - \pi_t\right)\left(\Gamma_t^0 + E_t \left(\eta_{t+1} \mid d = 0\right)\right). \tag{6}$$

If $E_t \left(\eta_{t+1} \mid d = 0\right) \simeq E_t \left(\eta_{t+1} \mid d = 1\right)$ and π_t is close to unity then $\Gamma_t^0 >> |\Gamma_t^1|$ so that the job loss shock is much greater than the job retention shock.

We expect that the effect of the job loss shock is heterogeneous across workers. We shall sometimes parameterize the job loss shock in terms of observables at time t. Specifically:

$$\Gamma_t^0 = \gamma_t^0 + \gamma' \mathbf{z}_t \tag{7}$$

where z_t is a vector of observable factors that affect the size of the job loss shock. These include tenure in the current job, gender, union status, and age. Of course, many determinants of the job loss shock may be unobservable, such as the job match quality or family financial circumstances; these are captured by γ^0.

The analysis above focuses on the unobservable *mue*. The next step is to relate this to (observable) consumption. Denote consumption in period t by c_t. We take the following form for consumption growth:

$$\Delta \ln c_{t+1} = \Delta\phi_{t+1} - \Delta\lambda_{t+1} \tag{8}$$

where the time varying factor $\Delta\phi_{t+1}$ includes anticipated changes in factors that affect utility (for example, age, marital status or children). This includes potentially observable factors and unobserved factors.[2] Substituting in (1) and (2) and taking expectations

[2] While we think of this formulation as an approximation, it is worth noting that it holds exactly if (i) the agent has a rate of time preference equal to the interest rate; and (ii) the agent has the following

conditional on displacement, $d = 0$, we have:

$$E_t\left(\Delta \ln c_{t+1} \mid d = 0\right) = \Delta\phi_{t+1} - \Gamma_t^0 - E_t\left(\eta_{t+1} \mid d = 0\right). \tag{9}$$

To simplify notation, we henceforth denote $E_t\left(\eta_{t+1} \mid d = 0\right)$ by μ_t^0 (and $E_t\left(\eta_{t+1} \mid d = 1\right)$ by μ_t^1) so that we can express a households expected consumption growth conditional on job loss as:

$$E_t\left(\Delta \ln c_{t+1} \mid d = 0\right) = \Delta\phi_{t+1} - \Gamma_t^0 - \mu_t^0. \tag{10}$$

The objects we aim to estimate are the average effect of job loss on the *mue*, on those who experience job loss (denoted $E^H[\Gamma_t^0 \mid d = 0]$) and the relationship between that mean and observable characteristics, $\mathbf{z}_t$. Note that $E^H[]$ denotes an average *across* the population of job losers and H indexes job losers. Thus $E^H[\Gamma_t^0 \mid d = 0]$ is the analogue, in this context, of the "average effect of the treatment on the treated", where the "treatment" is job loss.[3] Given the parameterization in equation (7), the relationship between this quantity and characteristics ($\mathbf{z}_t$) is captured by γ. Equation (10) says that the expected consumption growth for job loser is the sum of anticipated changes, the effect of the job loss and the effect of other shocks, given that the agent is displaced.

per period utility function:

$$\begin{aligned} u(c_t) &= (\phi_t - 1 - \ln c_t)c_t; \\ \phi_t - 1 - \ln c_t &> 0. \end{aligned}$$

This utility function has the usual properties: positive marginal utility, risk aversion (a negative second derivative) and prudence (a positive 3rd derviative.)

[3] Note that this means that we are not estimating the average effect of job loss on a worker selected at random from the pool of workers that could potential lose a job, but rather the average effect on actual job losers.

Averaging across job losers gives:

$$
\begin{aligned}
& E^H \left[E_t \left(\Delta \ln c_{t+1} \mid d=0 \right) \mid d=0 \right] \\
= \;& E^H[\Delta\phi_{t+1,h} \mid d=0] - E^H[\Gamma^0_{t,h}(\mathbf{z}_{t,h}) \mid d=0] - E^H[\mu^0_{t,h} \mid d=0] \\
= \;& E^H[\Delta\phi_{t+1,h} \mid d=0] - E^H[\gamma^0_{t,h} \mid d=0] \qquad (11) \\
-E^H[\boldsymbol{\gamma}'\mathbf{z}_{t,h} \quad \mid \quad & d=0] - E^H[\mu^0_{t,h} \mid d=0].
\end{aligned}
$$

If we could assume that $\mu^0_{t,h}$ and $\Delta\phi_{t+1,h}$ are uncorrelated with observed job loss cost factors ($\mathbf{z}_t$) then we could simply regress consumption growth on the observables, for the sample of displaced workers, to estimate γ. This is unsatisfactory in two respects. First, $E^H[\gamma^0 \mid d=0]$ is not identified separately from $E^H[\Delta\phi_{t+1,h} \mid d=0]$ and $E^H[\mu^0_{t,h} \mid d=0]$. This means that, even with this assumption, we can only estimate how the cost of job loss varies with observable characteristics $(\mathbf{z}_{t,h})$ but not the overall level of effect. The overall level is crucial from a policy point of view, where we may (for example) wish to relate the average costs of job loss to public expenditures on a proposed labour market program.

Second, the assumption that $\mu^0_{t,h}$ and $\Delta\phi_{t+1}$ are uncorrelated with observed job loss cost factors $(\mathbf{z}_{t,h})$ is difficult to maintain. For example, it would require that the effect of all other shocks $\left(\mu^0_{t,h}\right)$ not vary with the observed determinants of the job loss shock such as age, occupation or education. As regards to $\Delta\phi_{t+1,h}$, this includes life-cycle factors that are also likely correlated with γ^0 and $\mathbf{z}$: for example, age, material status, and family type. If $\mu^0_{t,h}$ and $\Delta\phi_{t+1}$ are correlated with observed job loss cost factors $(\mathbf{z}_{t,h})$ then the regression of consumption growth on the observables, for the sample of displaced workers, does not even identify the way that the cost of job loss varies with observable characteristics.

Thus plausible identification of $E^H[\Gamma_t^0 \mid d=0]$ and $\boldsymbol{\gamma}$ requires that we have a way to estimate $E^H[\Delta\phi_{t+1,h} \mid d=0] + E^H[\mu_{t,h}^0 \mid d=0]$. Our strategy for doing so is to use standard matching methods to construct a control group. This exploits the additivity assumed above plus the usual kind of conditional mean independence assumptions:

$$E^H\left(\mu_{t,h}^0 \mid d=0, X\right) = E^H\left(\mu_{t,h} \mid \mathbf{X}\right) \tag{12}$$

$$E^H\left(\Delta\phi_{t+1,h} \mid d=0, X\right) = E^H\left(\Delta\phi_{t+1,h} \mid \mathbf{X}\right) \tag{13}$$

where $\mathbf{X}$ is a set of observable characteristics used to match treatments and controls. In addition, we will require that standard common support conditions are satisfied.[4]

A key point of this paper is to suggest that controls drawn from workers experiencing continuing employment are unlikely to be appropriate, for two reasons. First, those in continuing employment may be sufficiently different from job losers that it may not be possible to adjust for the differences between them on the basis of observables. Second, and more subtly, expected consumption growth for someone who is not displaced is the sum of anticipated changes, the effect of other shocks and the (positive) effect of job retention:

$$E_t\left(\Delta \ln c_{t+1} \mid d=1\right) = \Delta\phi_{t+1} - \Gamma_t^1 - \mu_t^1 \tag{14}$$

Again, among workers with imperfect insurance against job losses, job retention is a favorable resolution of uncertainty and this positive economic shock results in a welfare gain. Thus even if the necessary conditional mean independence holds, we cannot use those in continuing employment to estimate $E^H[\Delta\phi_{t+1,h} \mid d=0]$ and $E^H[\mu_{t,h}^0 \mid d=0]$

[4]Before doing this, we will present crude estimates based on simple differences in mean consumption growth.

because their consumption growth is confounded by Γ_t^1.[5]

The solution we propose is to draw our controls from workers who experience temporary layoffs with a definite recall date. First, these workers may be more similar to job losers, and hence make an estimation strategy based on correcting for observable differences more palatable. More importantly, these workers are receiving insurance against job loss from their firms, so that plausibly:

$$\Gamma_t^0 \approx \Gamma_t^1 \approx 0 \tag{15}$$

for these workers. While temporary layoffs involve an income loss, it is transitory, with no loss of job match or firm specific human capital. Workers on temporary layoff are eligible for unemployment insurance, and in some cases temporary layoff procedures are carefully integrated with unemployment insurance provisions (for example, some workers temporarily laid off from unionized firms receive a firm or union funded top up to their unemployment insurance benefits.)[6] Thus the consumption growth of workers experiencing temporary layoffs can be used to estimate $E^H[\Delta\phi_{t+1} \mid d = 0]$ and $E^H[\mu_{t,h}^0 \mid d = 0]$.

To summarize, the empirical strategy that is motivated by the theoretical considerations above, and which we will implement in this paper is as follows. To estimate the

[5] How large might this bias be? Manski and Straub (2000) report that in the mid - 1990s, a sample of American workers had an subjective expecation of job loss of 15%. If $\mu_t^0 = \mu_t^1 = 0$, Equation (6) implies that:

$$\Gamma_t^1 = -\frac{(1-\pi)}{\pi}\Gamma_t^0 \approx 0.18\Gamma_t^0.$$

Note, however, that our sample of permanent job losers very likely had higher than average probabilities of job loss. If we selected continuously employed controls to match these job losers on the basis of observable characteristics, it is very likely that we also select controls who had a higher than average *ex ante* expectation of job loss. This means a larger π, and hence a larger bias.

[6] Of course, the temporary layoff may reveal information about the future viability of the firm - and hence its ability to continue to provide insurance against job loss in the future. Alternatively, for a firm whose continued operation is in doubt, a temporary layoff may be a *positive* shock. Our claim is only that such considerations are second order, so that temporary layoffs provide a good first approximation to a full insurance benchmark.

cost of job loss:

1. Use consumption growth to measure innovations in the *mue.*

2. Among the "treatment" group of job losers, consumption growth confounds the effects of job loss with the effects of other shocks and anticipated changes in the *mue*. These confounders would affect consumption growth under the counterfactual of full insurance against job loss shocks.

3. Construct a matched control group drawn from workers experiencing temporary layoff. Use this group to estimate consumption growth under the counterfactual of full insurance against job loss shocks.

4. The difference in consumption growth between the job losers and matched controls is an estimate of the cost of job loss among the job losers.

We now turn to a description of the data on which we implement this strategy.

3 Data

3.1 Survey

The data for this paper are drawn from a panel survey on Canadians who separated from a job: the Canadian Out of Employment Panel (COEP). The survey was conducted by Human Resources Development Canada (HRDC) to evaluate the effects of a series of changes in the Canadian Unemployment system in the mid- 1990s. A sample of some 11,000 workers who had a job separation in February or May of 1993 were interviewed three times, at about 26, 39 and 60 weeks after the job separation. In Canada, when a job separation occurs, the employer is obliged to file a "Record of Employment" (ROE) with

HRDC. These reports are compiled into the database from which the sampling frame was constructed. The sample is representative of job separations in relevant window. We refer to the job separation that led to inclusion in the sample as the "reference" separation.[7] Interviews were conducted over the telephone and took an average of 25 minutes.

A second sample of some 8,000 individuals who separated from a job in February or May of 1995 was subsequently drawn (this sample is again representative of separations in the relevant window). The survey instrument was refined (and slightly expanded) for this second survey but care was taken to insure backwards comparability. In addition, the third interview was dropped. Together, the 1993 and 1995 COEP surveys provide a large sample of individuals who separated from a job. The period of 1993 to 1995 was one of slowly improving labour market conditions in Canada (for example, the aggregate unemployment rate fell from 11.2 to 9.5%).

A feature of the data is the wide range of questions were asked including questions on the pre-separation job and reason for separation; labour market activity; job search details; the activities of other household members; income; expenditure and assets. The availability of expenditure data in a survey of this type is somewhat unique; further details on these questions are given below.

In this paper our primary focus is on information about expenditures in the period prior to the job separation (collected retrospectively at the first interview) and at the last opportunity we have to observe the respondents (the third interview for respondents in the 1993 sample and the second interview for respondents in the 1995 sample). The timing

[7] Because the administrative records that form the sampling frame are not complete until some months after the job separation, it was not possible to have the first interview closer to the separation date. Thus survey information about the periods just before and after the job separation are asked retrospectively from a point some 6 months on. This long interval between the job separation and the first interview is the price of a sample of only those who experience a job separation; this price is somewhat mitigated by the availability of complimentary administrative data which is collected continuously.

of the interviews was adjusted between the 1993 and 1995 samples so that the timing (relative to the job separation) of the third interview for the former sample corresponds roughly to the timing of the second interview for the latter sample. The details of interview timing are presented in Table A1, in the Appendix.

One reason to focus on the last point at which respondents are observed is that we wish to examine the change in the marginal utility of wealth ("permanent income") across a job loss. At earlier interviews, as smaller fraction of respondents are back in some employment and a greater fraction of the sample may be liquidity constrained. Where respondents are liquidity constrained our analysis of the permanent shock is confounded.[8]

3.2 Sample

With regard to sample selection we begin considering only respondents between the ages of 20 and 60, and exclude single adults living with parents or unrelated adults. Extensive experience with the data (as well as common sense) suggests that the latter group return expenditure information which is of poor quality. We also exclude workers who held multiple jobs at the separation date, one of which was ongoing.

Next we limit the sample to workers whose "reference" job had a duration of 6 months or more. This corresponds to the notion that a job loss presumes some attachment to the job. In fact, many studies have defined displaced workers as having "established work histories" (Kletzer, 1998) and some studies have limited their analysis to workers who lost jobs in which they had rather considerable tenure (for example, Jacobson, Lalonde and Sullivan (1993)). In our empirical analysis differences across workers with different

[8]The 1995 data contain direct questions about credit constraints between the job loss and the first interview, and at the first interview data. These have been analyzed by Crossley and Low, (2004). The reported incidence of binding credit constraints in these data is quite low. About a one in four permanent layoffs report being unable to borrow at the first interview, but only about one in twenty-five report that this is a binding constraint.

levels of pre-separation tenure will be an important focus.

We use self reported (survey) information to identify layoffs and quits, and to exclude other separation types (such as retirement).[9] We then limit the quit group to those who self reported that they quit to take another job. We have 402 such individuals. While layoffs are our primary focus, these voluntary job switchers provide some useful contrasts. In particular, among workers who voluntarily moved to (presumably better) jobs, one would expect that the shock of the job separation is, if anything, positive (the *mue* falls).

Among the layoffs, we distinguish types of layoffs on the basis of a series of survey questions about the *ex ante* (at time of layoff) expectation of recall. We define workers to have had a strong expectation of recall if they expected to be recalled on a specific date. We also refer to this group as "temporary layoffs". Those workers who reported no expectation of recall are our "permanent layoffs" and this is the principal group of interest for this study. Note that this *ex ante* definition of job loss or "displacement" differs from much of the displaced worker literature in which "displacement" is defined in terms of *ex post* realizations. However, conditioning on "time 0" information is much more natural in the consumption growth framework developed in the previous section. We also have a group of workers who expected recall but reported that they did not have a particular date by which they expected to be recalled. We refer to these workers as having "some expectation of recall".

Our data contain 3028 "permanent layoffs" (no expectation of recall), 1094 "temporary layoffs" (strong expectation of recall) and 1419 workers with some expectation of recall. The large number of temporary layoffs may be surprising to readers from outside

[9] The data also contain an administrative reason for separation (from the ROE form). These correlate reasonably well with self reported reasons, but have the drawback of a very large "other" category.

North America, but the important role of temporary layoffs in unemployment in North American labour markets is well documented (see for example, Feldstein, 1976).[10]

Table A2 in the Appendix documents the demographic and economic characteristics of respondents in each of the four groups just defined. The first panel of Table A2 reports demographic characteristics. The most dramatic differences - in terms of age, education, and local labour market conditions - are between quits and layoffs. The second panel of this Table reports economic characteristics prior to the reference separation. Relative to all layoffs, the quits have much shorter tenures on average. Comparing the temporary and permanent layoffs we note that the temporary layoffs are more likely to be unionized and have higher tenures. Note also that more than 80% of them expected the layoff. This further supports the notion that for this group, the shock associated with actual separation may be small, and thus that they may provide a good approximation to the full insurance benchmark.

In Table A3 we document the employment outcomes for these groups as of the first interview. There is attrition in our sample between the first and last interviews (see the first few rows of Table A2). In Table A3 we report the same first interview information for all first interview respondents (in the top panel) and for the sub sample that subsequently responded to the second interview (in the bottom panel). Comparing the top and bottom panels we note that the numbers are very similar. Thus this very simple exercise does not reveal any evidence that the attrition was nonrandom.

In terms of the actual outcomes we note that re-employment is much higher among temporary layoffs and quits than permanent layoffs. A small number of *ex ante* permanent layoffs do return to their former firm, while some *ex ante* temporary layoffs take work

[10] It's worth noting that the Canadian Unemployment Insurance system (unlike the U.S. system) has no experience rating of firms.

else where. If not re-employed, a permanent layoff is more likely to be actively searching than a temporary layoff. Workers with "some expectation of recall" exhibit outcomes which lie somewhere between the permanent and temporary layoff groups.

In Appendix Table A4 we summarize the labour market outcomes for these groups at the final interview. Interestingly, the employment rate among temporary layoffs fall from the first to final interview. This may be because the final interview is in the fifth quarter after the reference separation, and temporary layoffs are often seasonal in nature (even in non-seasonal manufacturing industries). By this point some 15% (26% of the 57% employed) of *ex ante* permanent layoffs have returned to their former firm, while have almost 20% (29% of the 66% employed) *ex ante* temporary layoffs are working at a new firm.

3.3 Expenditure Questions

For the purposes of this paper the most important set of variables are those concerning expenditures. Two sets of questions were asked at each interview. The first was a set of levels questions concerning expenditures in the past week or month on a range of goods including housing; food at home; food outside the home; clothing and total expenditures in a month. The second set comprised a single question regarding the change in total expenditures relative to the month prior to the ROE (separation) date. In this paper our focus in on total expenditures. This is consistent with the theoretical framework developed in the previous section. It is also the only (expenditure) quantity for which we have pre-separation information. Since these questions are somewhat unusual in a survey of this type, we present the full text of the questions here. At each interview, the respondent was asked:

About how much did you and your household spend on everything in the past month? Please think about all bills such as rent, mortgage loan payments, utility and other bills, as well as all expenses such as food, clothing, transportation, entertainment and any other expenses you and your household may have.

And:

Has the amount you spend on everything decreased since <ROE>?

By what amount monthly?

Has the amount you spend on everything increased since <ROE>?

By what amount monthly?

The first question provides c_{t+1} (consumption at the interview date) and the following sequence Δc_{t+1} (the change since just prior to the reference separation).We approximate $\Delta \ln c_{t+1}$by $\frac{\Delta c_{t+1}}{c_{t+1}}$.[11]

Although the answers to these questions are undoubtedly noisy, we have several reasons to believe that they contain significant information about the levels and changes in household expenditures. First, we note that in each survey the expenditure questions are asked before income questions, so that we think it is less likely that the respondents just report incomes in response to expenditure questions. Second in other work (Browning and Crossley, 2001, 2003; Browning, Crossley and Weber, 2004) and in unreported subsidiary analysis, we have amassed considerable internal and external evidence of the

[11]It is possible, of course, to construct c_t (consumption just prior to the reference separation) from c_{t+1} and Δc_{t+1},and then appoximate $\Delta \ln c_{t+1}$by $\frac{\Delta c_{t+1}}{c_t}$.However, for the relative small growth rates we consider, the two approximations differ little, and, in the present context and, it is likely that $\frac{\Delta c_{t+1}}{c_t}$would suffer from greater measurement error in the denominator.

validity of the expenditure responses in the COEP. In particular income elasticities and demographic effects can be precisely estimated with this data (which would not be the case if the data were simply noise) and the data perform well in a series of budget share and Engel curve comparisons with the FAMEX, a Canadian household budget survey thought to be of excellent quality.

3.4 A First Look at Earnings and Consumption Growth

Before turning to formal estimates of the costs of job loss, we provide a descriptive analysis of earnings and consumption growth from just before a job separation until the fifth quarter after job loss. Figure 1 presents box and whisker plots of proportional consumption and earnings changes for layoffs with strong expectation of recall (*i.e.*, a recall date), some expectation of recall, and no expectation of recall (permanent layoffs) as well as quits. In each case the left hand box reflects earnings growth and the right hand box consumption growth. A number of statistics corresponding to these pictures are presented in Table 1a. Differences across groups in earnings growth are stark. Five quarters out, the median individual who quit to take another job experienced substantial earnings growth (9%) while the median permanent layoff has earnings almost 50% below their pre-separation level. Both parametric tests of common means and nonparametric rank tests suggest that the distribution of proportional earnings changes of permanent layoffs is strongly statistically different from that of the other groups.

In contrast to earnings, the differences in consumption growth are not so visually striking. In every category the median change in consumption is zero. Nevertheless, those who quit to take another job do appear - in both the figure and in the mean - to experience stronger consumption growth than the other groups. The differences among the other groups are difficult to discern from the box and whisker plots, but the statistical tests reported in the bottom panel of Table 1a confirm that the permanent layoffs are different from each of the other groups. Temporary layoffs (strong expectation of recall) experience stronger consumption growth than those with some expectation of recall, who in turn experience more consumption growth than permanent layoffs (no expectation of recall). As noted in the introduction, there are a number of reasons to expect that any proportional change in individual earnings translates into a rather smaller change in household consumption (the earnings loss may be transitory, the individual may be providing only a fraction of household income). Nevertheless, the striking differences in earnings and consumption data, combined with the way the consumption data are collected, may suggest to some readers that the consumption data is simply noise. However, the statistically significant differences across groups, and the strong consumption growth of those who quit to take another job refutes that position.

As first reported in Table A1, the weeks elapsed between separation from the reference job and the final interview varies between approximately 54 and 64 weeks in our sample. The bottom row of Table 1a reports that the mean is between 58 and 59 weeks (about 9/8 of a year) for each of our separation type groups. Thus variation in elapsed time does not seem to have played any role in the heterogeneity in earnings and consumption growth across groups. Notice also that the data underlying both the figures and tables is nominal. This was a relatively low inflation period in Canada. The respondents to our sample experienced proportional changes in the CPI which ranged from -0.0018 to 0.027 (inflation of -0.1 to 2.7%). The bottom row of Table 1A reports that there was some difference in the inflation experienced across groups, with in particular the permanent layoffs experiencing on average one percentage point less inflation. This is a very small component of the differences in nominal consumption and earnings changes.

Figure 2 repeats the analysis of Figure 1, but with the sample limited to those who report being back in employment at the last interview. The corresponding statistics are reported in Table 1b. Several features of the Table bear notice. First, the differences in earnings growth across layoff groups largely disappear (in the means and figures - the rank tests still suggest statistically different distributions). Furthermore the median earnings change in each layoff group is non-negative. This suggests that among our sample the earnings changes associated with job separations are all associated with non-employment (and not with wage changes). This is inconsistent with studies of job displacement which have focussed on highly attached workers (for example Ruhm, 1991) which find that both wage and employment changes play a role, but it is consistent with studies such as Polsky (1999) which examine job losers of a broad range of labour force attachment. However, further breakdowns by tenure in the reference job revealed that in our data, as in most other studies, high tenure workers experience wage losses on re-employment.[12]

A key result of this analysis is the very strong consumption growth exhibited by voluntary job switchers, which averages 10% (over a period just longer than a year). A reasonable interpretation of the data is that these workers have experienced a significant positive shock to their lifetime wealth. This observation supports our assertion that great care must be taken in comparing displaced workers to workers who are not displaced.

We now turn to implementing the estimation strategy developed in Section 2. In doing so, we set aside the data on quits and on respondents with some expectation of recall, and focus on the permanent layoffs (our "treatment" group) and the temporary layoffs (those with a strong expectation of recall, from which we draw our controls.)

[12] Full results are available from the authors.

4 Estimation Results

The first row of Table 3a reports the average consumption growth for our full sample of permanent layoffs, and our full sample of temporary layoffs. The former report an average consumption loss of 3.1%, while the latter report average consumption growth of 2.5%. The difference between these averages is -5.6%.[13]

However, the permanent and temporary layoffs differ in many ways. Therefore we use a matching procedure to estimate the counterfactual consumption growth of the permanent layoffs. As noted above, our estimates then rely on a conditional independence assumption. We have two advantages in this regard. First, our data are quite rich, so it is possible to match treatments and controls on a wide range of observable pre-treatment (that is, pre-displacement) characteristics. Second, a substantial literature on the determinants of consumption growth provides guidance as to some of the factors that are important to control for. For example, we know that consumption growth varies significantly with age and family composition. Differences in these factors between the treatment and control could certainly lead to biased estimates of the counterfactual consumption growth of the treatment group, and so it is important to match on these characteristics.

To reduce the dimension of the matching problem, we match on the estimated propensity score. This was first suggested by Rosenbaum and Rubin (1983) and is now quite common in the evaluation literature (see, for example, Smith and Todd, 2003). In our context, the propensity score is the conditional probability that a worker in our sample

[13] Sample sizes reported here differ from those in our informal data analysis because of item non-response - either to the one or more of the consumption questions or to one or more of the co-variates we use to match treatments and controls. Because our data are rich and we control for a very large number of observable characteristics in our matching estimators, some loss of sample is inevitable.

is permanently (not temporarily) laid off.[14] The propensity score was estimated using a Probit model. The full set of explanatory variables in this model included a quadratic in age, gender, education dummies and the logarithm of household size; dummies for marital status and spousal employment status; dummies indicating capital income and home ownership; occupation dummies, a union dummy and job tenure dummies; a dummy for unemployment insurance use in the previous two years; a polynomial in earnings in the reference job; the local unemployment rate; broader region dummies and time dummies. The estimation results are presented in Table 2. The model has significant ability to discriminate between treatments and controls, with a 79% rate of correct classification. The pseudo R-square is 25%. The distributions of estimated propensity scores in the treatment and control groups are presented in Figure 3. We conducted a standard 'balancing test' and found that the balancing property was satisfied.[15]

With the estimated propensity score in hand, common support was imposed. This involved discarding 12 permanent layoffs with propensity scores greater than the largest propensity score among the temporary layoffs. Matching was then done by locally linear regression. The result of this exercise is reported in the second row of Table 3a. The average consumption group for permanent layoffs satisfying the common support condition is -3.0%. The average counterfactual consumption growth for this group (based on the matched controls) is 3.4%. Therefore, we estimate that this group experienced a consumption loss of 6.4%, relative to a benchmark of full insurance against job loss. This

[14] Propensity score matching estimators were implemented using PSMATCH2 in STATA. PSMATCH2 is generously made available by Leuven and Sianesi (2003).

[15] This test involves testing for mean differences in conditioning variables between the treatments and controls within strata of the propensity scores. It was implemented with PSCORE in STATA. PSCORE is generously made available by Becker and Ichino (2002).

Because this test involves many comparisons (26 conditioning variables by 8 strata of the estimated propensity score) we followed the advice of Lee (2006) and made a Bonferroni approximation to maintain the size of the test at 5%.

is the Average Effect of the Treatment on the Treated. A 95% confidence interval for this estimate was constructed using the percentile method on 999 bootstrap replications. This confidence interval, which accounts for the fact that the propensity score is estimated, is -9.9% to -3.6%. Thus the average consumption loss is certainly statistically different from zero.

We also investigated the robustness of this estimate to different aspects of the matching procedure. In particular, we (i) halved and doubled the bandwidth for the locally linear regression; (ii) trimmed the 5% of treatments whose propensities scores corresponded to the lowest estimated densities among controls; (iii) matched on the index rather than predicted probability; and (iv) used a single nearest neighbor match rather than locally linear regression. The resulting point estimates, which are reported in Appendix Table A5, ranged from -5.5% to -6.7%, indicating that our baseline estimate is robust to these choices.

Matching on a rich set of covariates involves some loss of sample because of item nonreponse. One of the lessons drawn by Smith and Todd (2005) is that results can be sensitive to changes in sample. To check this, we repeated the procedure with a more parsimonious specification of the propensity score, and consequently a larger sample of respondents. This produced a point estimate of -5.9%, which is reported in Appendix Table A5. This further increases our confidence in the robustness of our baseline estimate.

As suggested in Section 2, it is of interest to know not just the average loss with displacement, but how that loss varies with observable characteristics. Accordingly, we repeated the above exercise for different subsamples of the permanent layoffs. The results are reported in Table 3b. Rows 1, 2 and 3 deal respectively with permanent layoffs (and matching controls) who lost a union job, who were over 40 years of age, and who lost a job in which they had 10 or more years of tenure. (These groups are obviously not mutually exclusive). The results indicate that older workers, and workers with high job tenure experience particularly large consumption losses if permanently laid off. The point estimates are -9.3% and -10.4% respectively. In Row 4 of Table 3b we focus on female treatments and controls, where we find an estimated average effect of job loss of -6.9%.

Finally, in Table 3c, we attempt to deal empirically with two possible shortcomings of the theoretical framework developed in Section 2. The first issue we consider is that our theoretical framework associates the cost of job loss with the revision to the *mue* between the separation date (t) and a date just over a year later ($t + 1$). If displacement was known with certainty prior to our pre-separation consumption observation, then the theory suggests that information should be fully incorporated into λ_t and hence c_t. Our empirical strategy would then fail to capture the cost of job loss.[16]

We cannot examine consumption changes prior to the separation date with our short panel. However, respondents to the COEP surveys were asked (retrospective) questions about formal notice of lay off and whether they "expected" the lay off. Respondents were asked whether they received notice, and if so, how much notice they received. They were also asked whether they expected the lay off, and for how long they held this expectation. There is obviously some ambiguity regarding the interpretation of the latter. The pre-separation consumption question refers to the "month before job ended." To identify workers who clearly had information about the layoff prior to this period, we constructed dummy variables for receiving advance notice 6 or more weeks prior to the layoff, and for expecting the layoff 6 or more weeks prior to the layoff.[17] Unfortunately, in the 1993

[16] Existing evidence on this point is somewhat mixed. For example, Stephens (2001), using the P.S.I.D. reports that household food consumption falls prior to a job loss, presumably as the probability of job loss rises. However, Stephens (2004) reports that food consumption drops with job loss in the Health and Retirement Survey do not appear to vary with households subjective job loss expectations.

[17] The choice of 6 weeks here reflects in part an ambiguity in the wording of the consumption question.

survey the questions regarding length of notice and length of expectation were asked of a random 20 percent of respondents. Thus using these questions results in a substantial loss of sample size (albeit one where the data are missing at random.) Reports of significant notice (or expectation) turn out to be uncommon in these data. Among the permanent layoffs for which we have complete information, 12 percent had formal notice of the layoff 6 or more weeks prior to the layoff date, while 24 percent reported expecting the layoff 6 or more weeks prior to the layoff date. In the first two rows of Table 3c we report the matching estimate of the effect of job loss on consumption growth while deleting these respondents. Deleting (from both treatment and control groups) those reporting formal notice of the layoff 6 or more weeks prior to the layoff date essentially leaves our point estimate unchanged, at -6.6%. Deleting those reporting expectation of the layoff 6 or more weeks prior to the layoff date leads a point estimate of -10.4%. This is somewhat larger than our baseline estimate and can be interpreted as the average effect of an *unexpected* displacement on those who experienced an unexpected displacement.

The second issue we consider is that our development in Section 2 (especially Equation (8)) assumes separability of consumption and labor supply. Without this separability, changes in consumption might reflect substitutions with leisure rather than changes in the *mue*. As a rough check on this possibility, we implemented our estimator on the subsample of layoffs that were in employment at the last interview (so that labor supply was broadly similar at the pre- and post-displacement observations.) The results of this exercise are reported in the last row of Table 3c. This sample restriction leads to a point estimate of -4.6%. This is slightly smaller (in absolute value) than our baseline estimate,

It is not clear wether respondents would interpret "the month before the job ended" as the 30 days terminating with the layoff date or the last full calendar month prior to the layoff date. In either case, repondents with notice (or expectation) 6 or more weeks before the layoff date certainly had this information prior to period for which they report pre-separation consumption.

but does not significantly change our assessment.

5 Policy Implications.

Workers who lose their job because of a plant closure or other permanent mass layoff face a potential loss of lifetime earnings and welfare. Such workers disproportionately bear the costs of the adjustments necessary for the continued economic prosperity of their society. There is clear scope for social gains from providing insurance against these losses. In the absence of private insurance, some government provision is potentially useful. The actual impact of government policies will depend on the seriousness of such long run losses and their distribution across the working population. The main contribution of the current paper is to provide a novel method of calculating the long run costs of a job loss that does not require specification of the underlying labour market nor long panels on individual workers.

Our analysis departs from the literature on the costs of job loss in two main ways. First, we show that under specific assumptions, changes in consumption provide a good measure of the long run cost of a job loss. The prior literature focusses on wages, the duration of unemployment, and earnings after a displacement. It has proven very difficult to provide a credible link from these outcomes to long run welfare. In contrast, to the extent that households are rational and forward looking, consumption changes reflect both the household's ability to adjust to the shock, through a variety of means, and the household's expectations of the long run effects of the job loss. Thus consumption changes provide a measure of the cost of job loss that is both comprehensive and long run. In addition to the theoretical attraction of using consumption changes, there is an empirical advantage in that we do not need a long time series of observation on each household; a short panel will do.

Our second departure from the current literature is our emphasis on defining the appropriate counterfactual for policy analysis. The counterfactual we propose is the ideal situation in which society provides full insurance against job loss shock. Job loss costs relative to this benchmark provide an upper bound on the benefit of *any* new cost-mitigating policy introduced in the current economic environment. Actually estimating the cost of job loss requires a control group, and the choice of control group is strongly influenced by the choice of counterfactual or benchmark. The past literature on displaced worker has either foregone a control group, or (implicitly or explicitly) drawn controls from the pool of non-displaced workers in the data at hand. We argue that, given our choice of counterfactual, workers in continuing employment are poor controls since such workers experience the positive shock of job retention. In our analysis, we employ an alternative and novel source of controls: workers experiencing temporary layoffs with a known date of recall. Firms that use temporary layoffs to manage demand shocks effectively insure their workers against job loss. Thus these workers provide a way to approximate average consumption growth under the counter factual of full insurance.

We apply our methods to a Canadian survey of workers who experience a job separation in the mid-1990's. We find that permanently displaced workers suffer a consumption loss of between 4% and 10% of pre-job-loss consumption, with a point estimate of 6.4%. One way to assess this number is to note that our estimate of consumption growth under the counterfactual of full insurance is 3.4% over a period just longer than a year. Thus a job loss is equivalent to losing two years of 'normal' consumption growth. This is a very substantial loss. Moreover, the loss is higher for particular groups. For older workers the estimated average loss is 9.3% and for workers who had been in their firm for over 10 years the estimated loss is 10.4%. The latter estimate is striking given that long tenure Canadian workers who are permanently displaced are eligible for tenure related redundancy payments which are designed to offset the costs we are concerned with. The implication is clearly that current rules regarding redundancy payments do not provide anything like adequate insurance. This is not to recommend, of course, that we should increase current levels of redundancy payments, since this may have other deleterious effects, but it does suggest that provisions in place in Canada in the mid-1990's failed to provide adequate insurance against a considerable individual risk. This leaves considerable scope for government policies to mitigate the long run impacts of job loss.

6 Conclusions.

We provide a consumption based method for measuring the long run impact of displacement from a job. Consumption losses are conceptually different from earnings losses. Earnings losses may be persistent, but both theory and data suggest that they are not fully persistent: there is eventually some catch up since most workers find new employment, and there may be some recovery of wages as they accumulate new firm-specific human capital or match quality. In contrast, the theoretical framework developed in Section 2 suggests that consumption losses are roughly permanent. The intuition behind this claim is that a forward looking household's post-displacement consumption choice already reflects their best guess of future earnings, including any anticipated catch up. The principal virtue of our Euler equation based approach is that it does not require specifying an underlying model of the labour market or the possibilities for self-insurance through occupational choice, saving behaviour, work adjustments by a spouse or existing governmental social insurance systems. They thus provide a relatively robust estimate of the costs of a job loss. They do not, however, provide a vehicle for the analysis of specific policy suggestions for providing workers with better insurance against long run losses. For that we do need structural estimates of the labour market and the insurance possibilities open to individual workers. Such an analysis would provide a framework for considering jointly consumption and labour supply and hence analyzing policy suggestions that would be partly designed to offer workers better insurance against the very substantial losses we have found in this paper.

Acknowledgements

We gratefully acknowledge the financial support of the Danish National Research Foundation through their grant to CAM, as well as the Social Sciences and Humanities Research Council of Canada. We received very helpful comments from two anonymous referees, from participants in the conference associated with this special issue, and from participants in a number of other conferences and seminars. All remaining errors are our own.

References

[1] Becker, Sacha O. and Andrea Ichino, (2002), Estimation of Average Treatment Effects Based on Propensity Scores. *The Stata Journal*, 2(4):358-377.

[2] Browning, Martin and Thomas F. Crossley, (2001), Unemployment Benefit Levels and Consumption Changes. *Journal of Public Economics*, 80(1)1-23.

[3] Browning, Martin and Thomas F. Crossley, (2003), Shocks, Stocks and Socks: Smoothing Consumption over a Temporary Income Loss. McMaster University, Department of Economics, Working Paper 2003-07.

[4] Browning, Martin, Thomas F. Crossley and Guglielmo Weber, (2003), Asking Consumption Questions in General Purpose Surveys. *Economic Journal*, 113:F540-F567.

[5] Cochrane, John. H., (1991), A Simple Test of Consumption Insurance. *Journal of Political Economy*, 99(5):957-76.

[6] Crossley, Thomas F., and Hamish Low, (2004), When Might Unemployment Insurance Matter? Credit Constraints and the Cost of Saving. McMaster University, Department of Economics, Working Paper 2004-04.

[7] Fallick, Bruce, (1996), A Review of the Recent Empirical Literature on Displaced Workers. *Industrial and Labor Relations Review*, 50(1):5-16.

[8] Feldstein, Martin, (1976), Temporary Layoffs in the Theory of Unemployment. *Journal of Political Economy*, 84(5):937-957.

[9] Gruber, Jonathan (1997), The Consumption Smoothing Benefit of Unemployment Insurance. *American Economic Review*, 87(1):192-205.

[10] Kletzer, Lori G., (1998), Job Displacement. *Journal of Economic Perspectives*, 12(1):115-136.

[11] Kuhn, P., (2002), Summary and Synthesis. in P. Kuhn, Ed., *Losing Work, Moving On: International Perspectives on Worker Displacement*, Kalamazoo, Michigan: W.E. Upjohn Institute for Employment Research.

[12] Kuhn, Peter and Arthur Sweetman, (1998), Wage Loss Following Displacement: The Role of Union Coverage. *Industrial and Labor Relations Review*, 51(3):384-400.

[13] Lee, Wang-Sheng, (2006), Propensity Score Matching and Variations on the Balancing Test. manuscript, Melbourne Institute of Applied Economic and Social Research.

[14] Leuven, Edwin, and Barbara Sianesi, (2003), PSMATCH2: A STATA module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing. http://ideas.repec.org/c/boc/bocode/s432001.html.

[15] Manski, Charles F. and John D. Straub, (2000), Worker Perceptions of Job Insecurity in the Mid-1990s: Evidence from the Survey of Economic Expectations. *Journal of Human Resources*, 35(3):447-79.

[16] Polsky, Daniel, (1999). Changing Consequences of Job Separation in the United States. *Industrial and Labor Relations Review*, 52(4):565-580.

[17] Rosenbaum, P.R. and D.B. Rubin, (1983), The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, 70(1), 41-55.

[18] Ruhm, Christopher J., (1991), Are Workers Permanently Scarred by Job Displacement. *American Economic Review*, 81(1):319-24.

[19] Smith, Jeffery and Petra Todd, (2005), Does Matching Overcome Lalonde's Critique of Non Experimental Estimators? *Journal of Econometrics*, 125(1-2):305-53.

[20] Stephens Jr., Melvin, (2001), The Long – Run Consumption Effects of Earnings Shocks. *The Review of Economics and Statistics*, 83(1):28-36.

[21] Stephens Jr., Melvin, (2004), Job Loss Expectations, Realizations, and Household Consumption Behavior. *The Review of Economics and Statistics*, 86(1):253-269.

TABLE 1a: Descriptive Statistics: Earnings and Expenditure Changes
Pre- reference separation to last interview
Proportional Changes in nominal monthly Amounts
All Final Interview Respondents

		Layoffs			Quit
		No Expectation of Recall	Some Expectation of Recall	Strong Expectation of Recall	
Earnings	**q1**	**-1**	**-1**	**-1**	**-0.40**
	q2	**-0.47**	**-0.19**	**0**	**0.09**
	q3	**0.016**	**0.025**	**0.025**	**0.04**
	mean	**-0.44**	**-0.39**	**-0.31**	**-0.013**
	Difference of mean from no expectation group, [t-stat]		**0.044** [1.9]	**0.13** [5.1]	**0.42** [11.1]
	Kruskal-Wallis rank test of common distribution with no expectation group: $\chi^2_{(1)}$ (p-value)		**8.6** (0.003)	**36.5** (<0.001)	**109.2** (<0.001)
Total Expenditure	**q1**	**0**	**0**	**0**	**0**
	q2	**0**	**0**	**0**	**0**
	q3	**0.051**	**0.044**	**0.063**	**0.11**
	mean	**-0.033**	**0.005**	**0.023**	**0.067**
	Difference of mean from no expectation group, [t-stat]		**0.038** [4.2]	**0.056** [5.7]	**0.099** [6.7]
	Kruskal-Wallis rank test of common distribution with no expectation group: $\chi^2_{(1)}$ (p-value)		**11.6** (<0.001)	**30.0** (<0.001)	**39.0** (<0.001)
CPI % change	**mean**	**0.6**	**1.4**	**1.3**	**1.5**
Weeks elapsed	**mean**	**59**	**58**	**58**	**59**

TABLE 1b: **Descriptive Statistics: Earnings and Expenditure Changes**
Pre- reference separation to last interview
Proportional changes in nominal monthly amounts
Respondent Employed at Last Interview Only

		Layoffs			Quit
		No Expectation of Recall	Some Expectation of Recall	Strong Expectation of Recall	
Earnings	q1	-0.25	-0.045	0.0062	0.025
	q2	0	0.025	0.025	0.19
	q3	0.20	0.097	0.10	0.45
	mean	0.032	0.033	0.071	0.25
	difference of mean from no expectation group, [t-stat]		0.00018 [0.0]	0.038 [1.8]	0.22 [8.0]
	Kruskal-Wallis rank test of common distribution with no expectation group: $\chi^2_{(1)}$ (p-value)		6.9 (0.008)	19.3 (<0.001)	63.4 (<0.001)
Total Expenditure	q1	0	0	0	0
	q2	0	0	0	0.017
	q3	0.063	0.056	0.067	0.14
	mean	0.003	0.030	0.042	0.086
	difference of mean from no expectation group, [t-stat]		0.027 [2.7]	0.039 [3.8]	0.083 [6.1]
	Kruskal-Wallis rank test of common distribution with no expectation group: $\chi^2_{(1)}$ (p-value)		2.9 (0.09)	12.1 (<0.001)	25.3 (<0.001)
CPI, % change	mean	0.6	1.4	1.3	1.5
Weeks elapsed	mean	59	58	59	59

TABLE 2: PROBIT ESTIMATION OF THE PROPENSITY SCORE		
Dependent variable = 1 if Permanent Layoff	(No expectation of recall)	
=0 if Temporary Layoff	(Expects recall on a specific date)	
	Marginal Effect	Standard Error of Marginal Effect
High school graduate	0.011	[0.026]
College graduate	0.022	[0.031]
Age	0.003	[0.012]
Squared	-0.006	[0.011]
Ln(household size)	-0.167	[0.036]**
Male	0.068	[0.026]**
Couple, spouse not employed	0.025	[0.026]
Single	-0.185	[0.061]**
Lone parent	-0.033	[0.043]
Some capital income, previous year	0.017	[0.022]
Household owns home	-0.016	[0.026]
Manager, pre-separation job	-0.124	[0.030]**
Blue Collar, pre-separation job	-0.190	[0.031]**
Unionized, pre-separation job	-0.136	[0.025]**
3 -10 years tenure, pre-separation job	-0.140	[0.025]**
> 10 years tenure, pre-separation job	-0.155	[0.034]**
UI use in 1 of previous 2 years	-0.145	[0.021]**
Ln (real earnings, pre-separation job)	0.123	[0.032]**
Squared	0.026	[0.023]
Cubed	-0.013	[0.015]
Atlantic	0.037	[0.036]
Quebec	0.031	[0.027]
Prairies	0.135	[0.026]**
British Columbia	0.024	[0.040]
Local unemployment rate	-0.134	[0.323]
Separation in 1995	-0.358	[0.021]**
Observations	2248	
Correctly Predicted	79%	
Pseudo R-square	0.25	
Notes: Standard errors in brackets, * significant at 5%; ** significant at 1%		

TABLE 3a: The Effect of Permanent Job Loss on Consumption Growth - Baseline Estimates				
	Sample Sizes	Mean Consumption Growth (%)		
	Treated Controls	Permanent Layoffs (Treated)	Temporary Layoffs With Recall Date (Controls)	Difference [95% C.I.]
Unmatched Comparison	1461 657	-3.1	2.5	-5.6 [-7.73, -3.55]
Matched Controls, Common Support	1449 657	-3.0	3.4	-6.4 [-9.6, -3.8]

Notes to the Matching Estimate (2nd Row):

1. **Treatments and controls were matched on the estimated propensity score, and common support was imposed. The propensity score was estimated with a Probit. Matching was done by locally linear regression. Additional details are in the text.**
2. **The difference reported in the far right column is an estimate of the Average Effect of the Treatment on the Treated (where the treatment is job loss.)**
3. **The confidence interval was constructed using the percentile method on 999 bootstrap replications.**

TABLE 3b: The Effect of Permanent Job Loss on Consumption Growth - Subsamples (Matched Controls, Common Support)

	Sample Size Treated Controls	Mean Consumption Growth (%) Permanent Layoffs (Treated)	Temporary Layoffs With Recall Date (Controls)	Difference [95% C.I.]
Unionized	386 305	-2.6	2.1	-4.7 [-8.4, -0.8]
Age > 40 years	579 264	-6.4	2.9	-9.3 [-14.3, -4.5]
Job Tenure > 10 years	218 172	-7.4	3.0	-10.4 [-17.6, -3.8]
Women	701 347	-2.6	4.4	-6.9 [-12.0, -2.0]

Notes:

1. **Treatments and controls were matched on the estimated propensity score, and common support was imposed. The propensity score was estimated with a Probit. Matching was done by locally linear regression. Additional details are in the text.**
2. **The difference reported in the far right column is an estimate of the Average Effect of the Treatment on the Treated (where the treatment is job loss.)**
3. **Confidence intervals were constructed using the percentile method on 999 bootstrap replications.**

TABLE 3c: The Effect of Permanent Job Loss on Consumption Growth – Subsamples II (Matched Controls, Common Support)

	Sample Size Treated Controls	Mean Consumption Growth (%) Permanent Layoffs (Treated)	 Temporary Layoffs With Recall Date (Controls)	Difference [95% C.I.]
Advanced Notice < 6 weeks (incl. 0)	503 392	-1.6	5.0	-6.6 [-13.3, -0.9]
Expected Job Loss < 6 weeks (incl. 0)	599 293	-2.4	8.0	-10.4 [-14.7, -2.9]
Employed at Last Interview	780 399	-0.5	5.1	-4.6 [-7.5, -1.8]

Notes:

1. **Treatments and controls were matched on the estimated propensity score, and common support was imposed. The propensity score was estimated with a Probit. Matching was done by locally linear regression. Additional details are in the text.**
2. **The difference reported in the far right column is an estimate of the Average Effect of the Treatment on the Treated (where the treatment is job loss.)**
3. **Confidence intervals were constructed using the percentile method on 999 bootstrap replications.**

Appendix: Further Data Description and Robustness Checks

TABLE A1: Interview Timing, 1993 and 1995 COEP
(Weeks since Reference Separation; Inter-quartile Range)

	1993 Cohort 1	1993 Cohort 2	1995 Cohort 1	1995 Cohort 2
Reference Job Separation	Feb. - Mar.	Apr.	Jan.-Mar.	Apr.-June
Interview1	27-29	24-25	36-40	33-38
Interview 2	40-43	37-40	60-63	54-57
Interview 3	61-64	55-59	X	X

TABLE A2: Descriptive Statistics: Pre - Reference Separation Information				
	Layoffs			Quits
	No Expectation of Recall	Some Expectation of Recall	Strong Expectation of Recall	
1st Interview Obs.	3023	1417	1094	402
COEP 1995 (%)	845 (28%)	1122 (79%)	794 (73%)	344 (86%)
Last Interview Obs. (%)	2199 (73%)	1127 (80%)	890 (81%)	315 (78%)

Characteristics of Respondent and Household				
High school graduate	0.37	0.42	0.44	0.42
College graduate	0.33	0.21	0.27	0.43
Age	38.0	37.8	39.0	32.7
Male	0.53	0.61	0.48	0.60
Couple, spouse not employed*	0.22	0.24	0.20	0.19
Single*	0.17	0.19	0.13	0.22
Lone parent*	0.089	0.050	0.072	0.041
Ln (household size)	0.94	0.95	1.03	0.89
Some capital income, previous year	0.39	0.32	0.38	0.28
Household owns home	0.62	0.66	0.71	0.60

* Omitted household type is couple, spouse employed.

TABLE A2: Descriptive Statistics: Pre - Reference Separation Information (Cont´d)

	Layoffs			Quits
	No Expectation of Recall	Some Expectation of Recall	Strong Expectation of Recall	
Characteristics of Reference Separation Job				
manager	0.28	0.18	0.28	0.30
blue collar	0.33	0.61	0.46	0.29
union	0.27	0.42	0.47	0.15
seasonal	0.10	0.28	0.33	0*
expected loss	0.45	0.71	0.81	1*
Job Tenure (Months)	65.2	80.4	89.7	44.5
Monthly Earnings	1.89	1.76	1.65	1.76
Program Use				
UI in at least 1 of past 2 years	0.55	0.80	0.74	0.40
Labour Market				
Atlantic	0.08	0.13	0.11	0.09
Quebec	0.27	0.40	0.31	0.22
prairies	0.15	0.12	0.08	0.19
BC	0.10	0.10	0.06	0.12
Local unemployment Rate	10.5%	10.6%	10.1%	9.2%

TABLE A3: Descriptive Statistics: First Interview Information

	Layoffs			Quit
	No Expectation of Recall	Some Expectation of Recall	Strong Expectation of Recall	
All First Interview Respondents				
Employed	0.44	0.60	0.80	0.79
Of Employed: Back at reference Employer	0.13	0.75	0.90	0.08
Job as good as reference job*	0.82	0.89	0.90	0.96
Of Non-Employed: Still in First UE Spell	0.77	0.53	0.49	0.26
Searched in Last 4 weeks	0.82	0.72	0.59	0.59
Participation Rate	0.84	0.85	0.89	0.89
Last Interview Respondents Only				
Employed	0.43	0.61	0.80	0.79
Of Employed: Back at reference Employer	0.12	0.76	0.90	0.08
Job as good as reference job*	0.83	0.89	0.90	0.96
Of Non-Employed: Still in First Spell	0.77	0.52	0.46	0.28
Searched in Last 4 weeks	0.81	0.72	0.61	0.56
Participation Rate	0.84	0.85	0.89	0.89

* Including those back at reference employer

TABLE A4: Descriptive Statistics: Last Interview Information

	Layoffs			Quit
	No Expectation of Recall	**Some Expectation of Recall**	**Strong Expectation of Recall**	
Employed	**0.57**	**0.59**	**0.66**	**0.80**
Of Employed: Back at reference employer	**0.26**	**0.58**	**0.71**	**0.31**
Job as good as reference job*	**0.79**	**0.84**	**0.88**	**0.94**
Of Non-Employed: Still in first UE spell	**0.44**	**0.27**	**0.20**	**0.25**
Searched in Last 4 weeks	**0.68**	**0.60**	**0.54**	**0.33**
Participation Rate	**0.83**	**0.71**	**0.74**	**0.82**

*** Including those back at reference employer**

TABLE A5: The Effect of Permanent Job Loss on Consumption Growth – Robustness Checks

	Sample Size	Mean Consumption Growth (%)		Difference
	Treated Controls	Permanent Layoffs (Treated)	Temporary Layoffs With Recall Date (Controls)	
Halved the bandwidth in local linear regression used in matching	1449 657	-3.0	3.7	-6.6
Doubled the bandwidth	1449 657	-3.0	3.1	-6.1
Deleted the 5% of treatments whose propensity scores corresponded to the lowest estimated densities among controls	1388 657	-2.9	2.8	-5.7
Matched in the index rather than the predicted probability	1449 657	-3.0	3.3	-6.3
Used a single nearest neighbour match rather than locally linear regression	1449 657	-3.0	2.5	-5.5
More parsimonious specification of the propensity score model	1688 723	-3.5	2.4	-5.9

Notes:

1. **The difference reported in the far right column is an estimate of the Average Effect of the Treatment on the Treated (where the treatment is job loss.)**

Figure 1: Proportional Income and Expenditure Changes

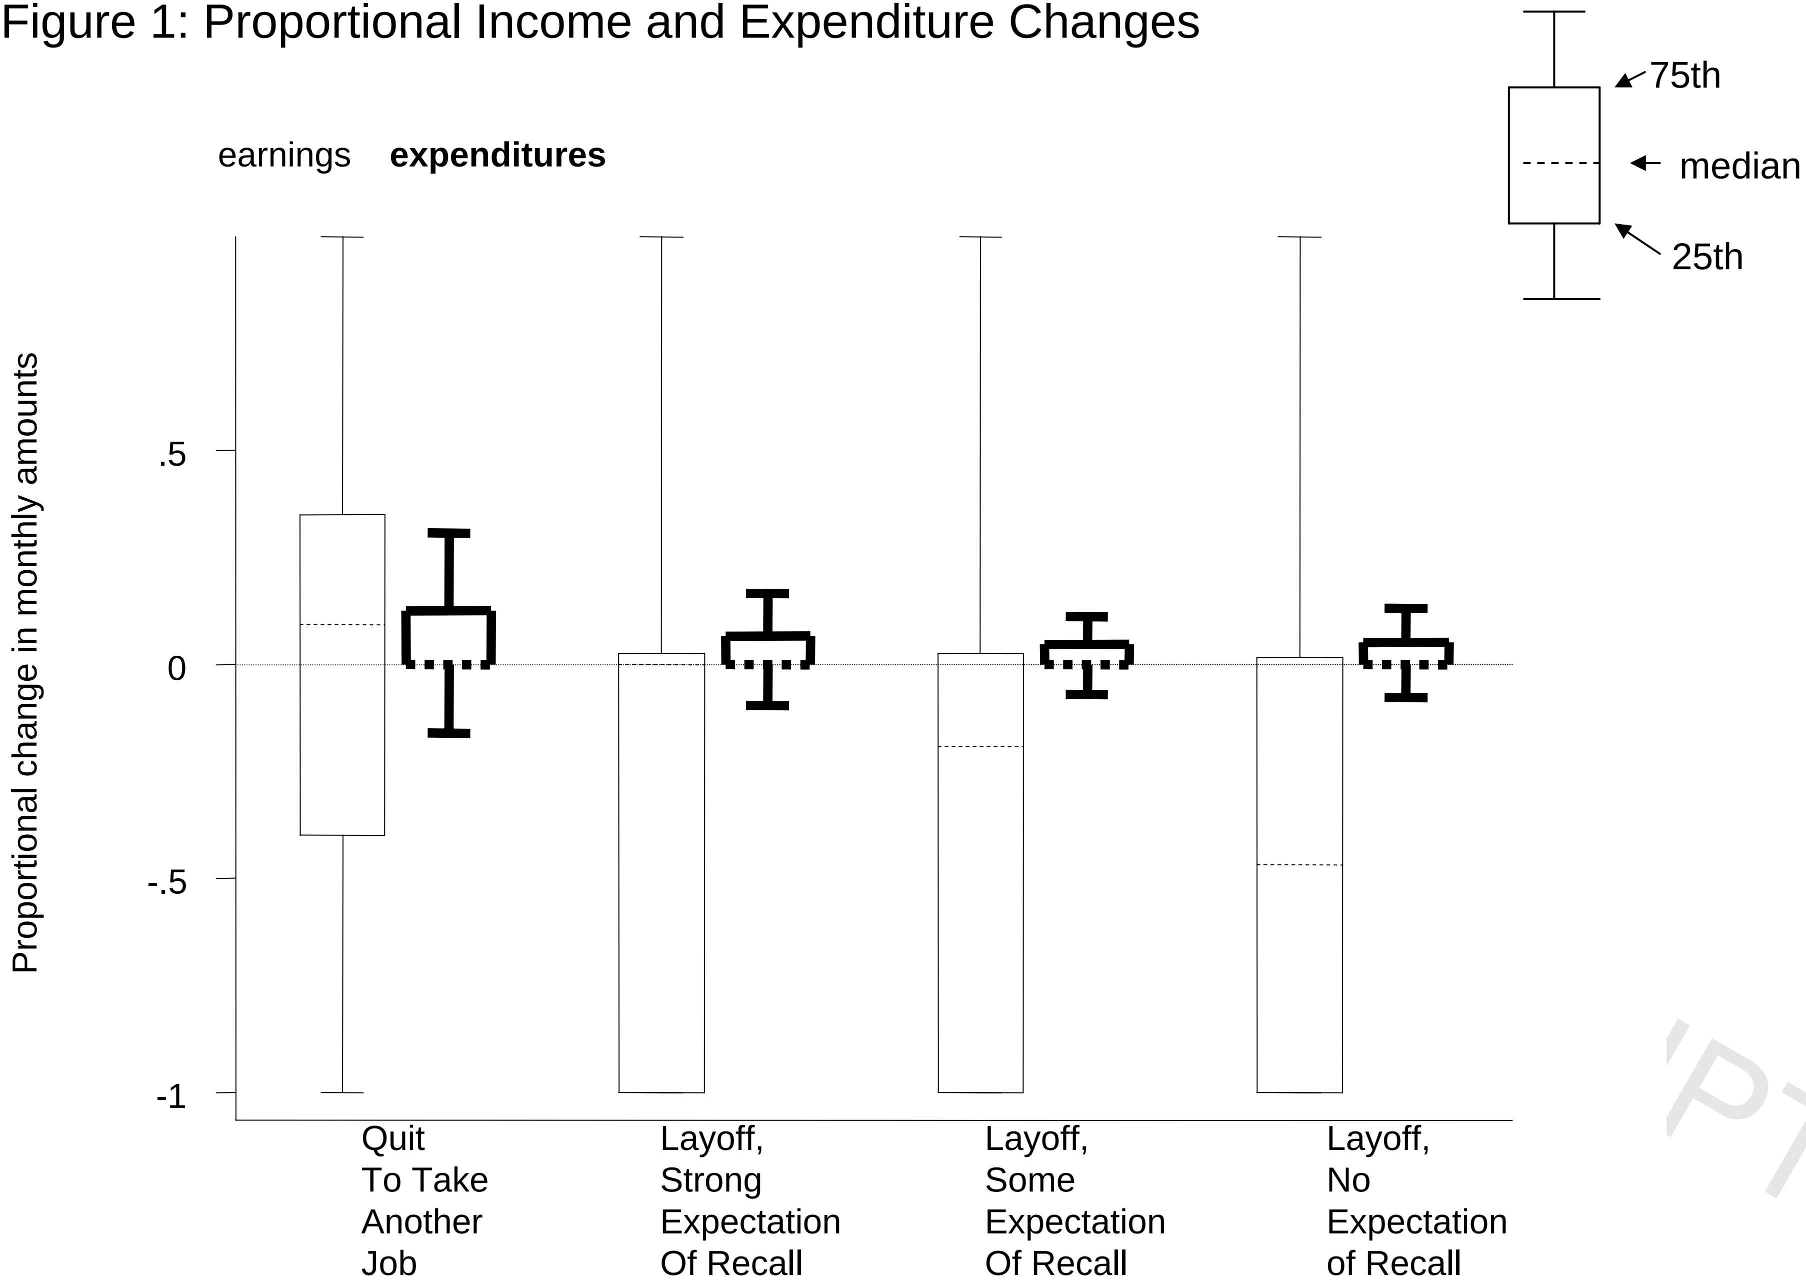

Figure 2: Proportional Income and Expenditure Changes

(Respondents Employed at Last Interview Only)

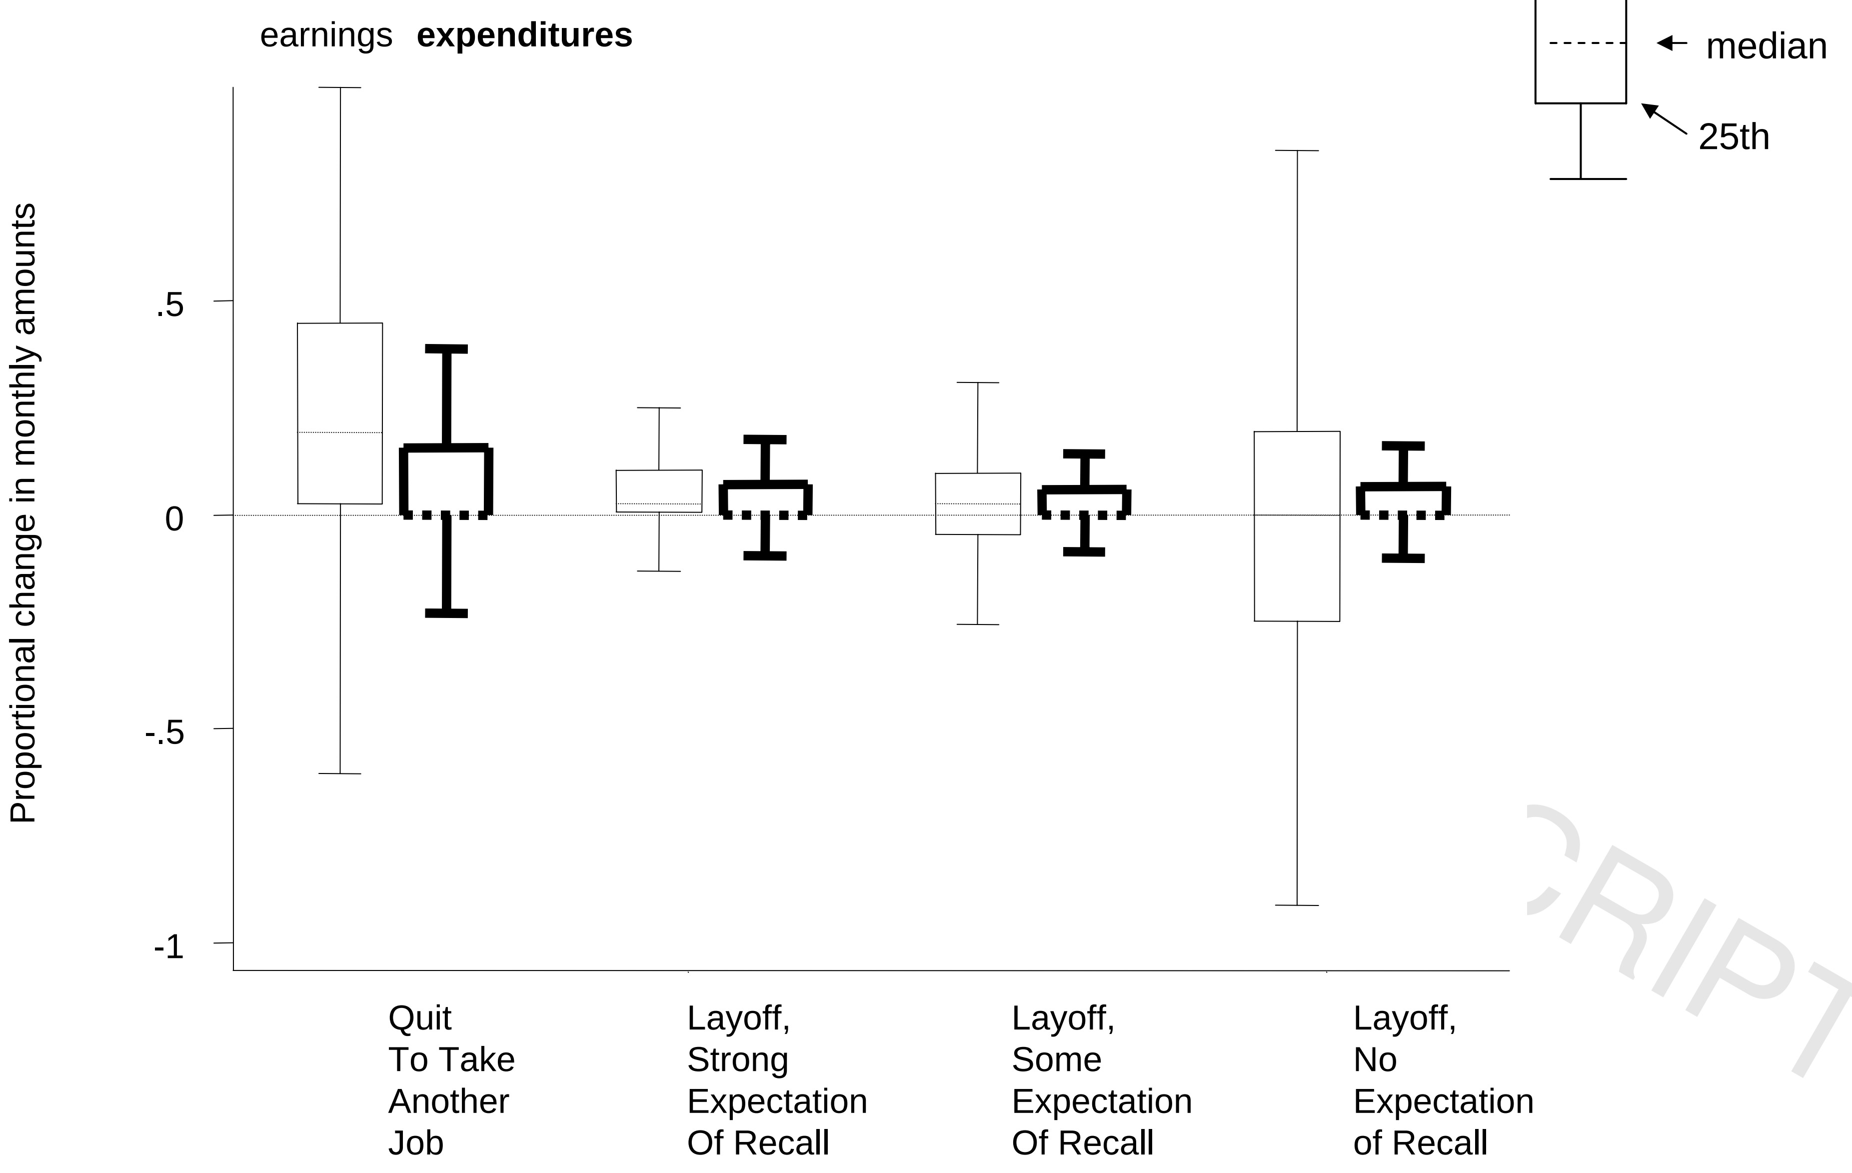

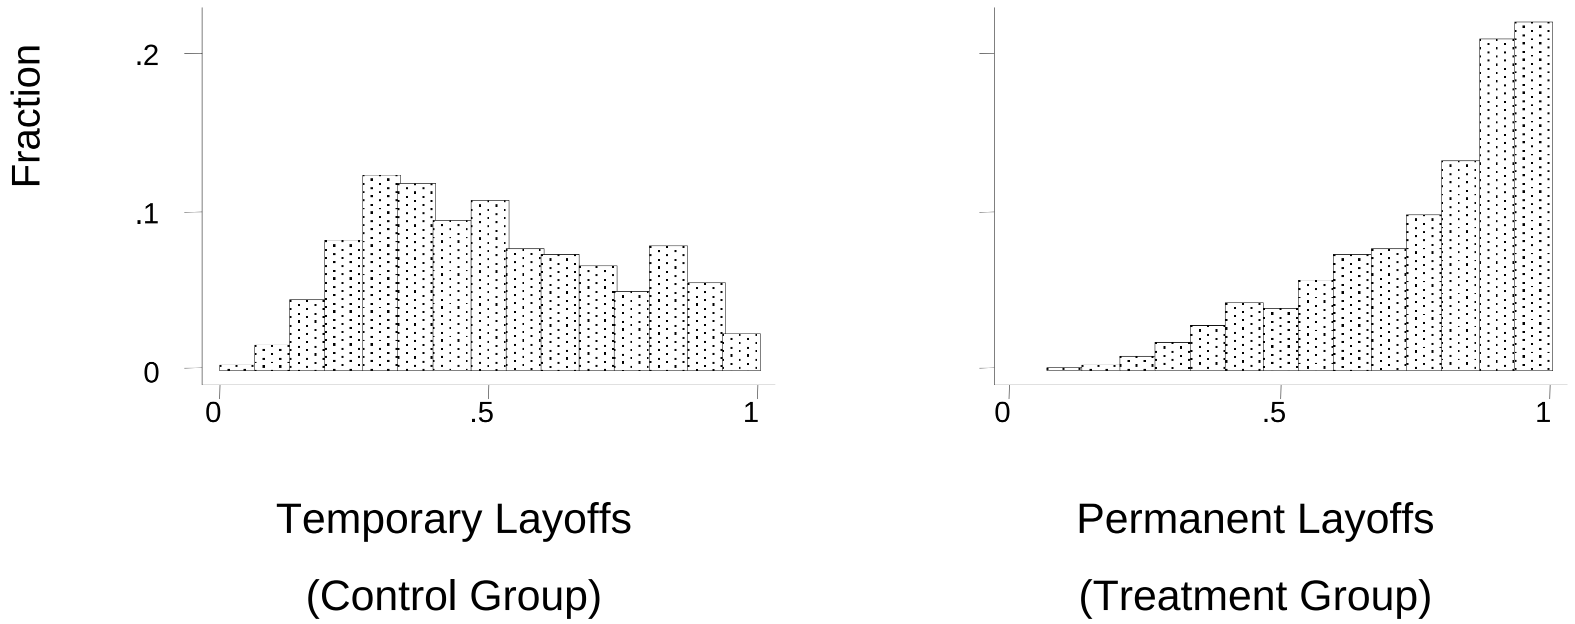

Figure 3: Distributions of Propensity Scores for Permanent Layoff