ROYAL STATISTICAL SOCIETY
DATA | EVIDENCE | DECISIONS

Journal of the Royal Statistical Society
Applied Statistics
Series C

Appl. Statist. (2019)
68, *Part* 3, *pp.* 793–807

A new competing risks decomposition: application to the effect of cutting unemployment benefit on unemployment durations

Donal O'Neill

Maynooth University, Republic of Ireland, and IZA Institute of Labor Economics, Bonn, Germany

[Received April 2018. Final revision November 2018]

Summary. I propose a competing risks decomposition of the difference in the restricted mean lifetimes of two groups, which distinguishes between the incidence and duration effects of exit states. I use the decomposition to evaluate the effect of a cut in unemployment benefits on unemployment durations in Ireland. Whereas the aggregate effect of the benefit cut is similar for men and women, the decomposition reveals that the channels through which the cut affects unemployment durations differ substantially.

Keywords: Competing risks; Cumulative incidence function; Decomposition; Unemployment duration

1. Introduction

In survival analysis, interest centres on the time until an event occurs. If several types of event can potentially occur, researchers and policy makers may wish to understand the effect of these competing risks. The typical approach that is used in social science estimates cause-specific competing risks hazard functions (Han and Hausman, 1990). However, this approach is not appropriate if we are interested in decomposing differences in durations in the time domain; although cause-specific hazards are useful when examining the rate of occurrence of an outcome, in the subset of people who are event free, they do not identify the absolute risk of a cause-specific outcome (Austin and Fine, 2017; Andersen *et al.*, 2012).

In this paper, I extend previous work on competing risks and propose a new decomposition of the difference in restricted mean lifetimes between two groups. In particular, I decompose this difference into the contributions of distinct exit states, distinguishing between duration and incidence effects of each state. The decomposition builds on the cumulative incidence function (CIF) approach that has previously been used in the biomedical literature and is valid irrespective of the shape of the underlying hazards or the dependence structure across exit states. The approach is easily implemented by using existing software, even in the presence of right-censored unemployment spells. I use the decomposition to re-examine the labour market response of 18-year-old jobseekers in Ireland to a 50% cut in unemployment benefits.

Address for correspondence: Donal O'Neill, Department of Economics, Finance and Accounting, Rhetoric House, Maynooth University, Maynooth, County Kildare, Republic of Ireland.
E-mail: donal.oneill@mu.ie

 0035–9254/19/68793

The program that was used to analyse the data can be obtained from

https://rss.onlinelibrary.wiley.com/hub/journal/14679876/series-c-datasets

2. Competing risks decomposition

A competing risk is an event whose occurrence precludes the occurrence of another event. One cannot observe people who exit unemployment to training, also ending the same spell with an exit to work. Identifying the importance of these alternative exit states is often crucial for policy makers. For instance, shorter unemployment durations, which are driven by exits from the labour market, are likely to be valued differently from shorter durations driven by exits to work.

When fitting models in the presence of competing risks, researchers can choose to model either the cause-specific hazard or on the cause-specific CIF. Cause-specific hazards estimate the rate of occurrence of a specific outcome, in the subset of the population who are event free at the specified time. This approach is widely used in the social science literature, e.g. when modelling unemployment duration (Katz and Meyer, 1990; Dolton and O'Neill, 1996; McVicar, 2008; van den Berg *et al.*, 2014). However, the cause-specific hazard does not provide the information that is needed for a competing risks decomposition of durations in the time domain. To see this, consider the simple case in which we are interested in decomposing a conditional expectation, $E(Y|Y \leqslant \tau)$. (Later we shall extend this to consider decompositions of the τ-restricted mean lifetimes, $E\{\min(Y,\tau)\}$, and show that similar issues arise.) Assuming K possible exits and using the law of iterated expectations, we can write this as

$$E(Y|Y \leqslant \tau) = \sum_{k=1}^{K} \frac{\Pr(K=k, Y \leqslant \tau)}{\Pr(Y \leqslant \tau)} E(Y|K=k,\ Y \leqslant \tau).$$

A key component in this expression is the cause-specific CIF, $\mathrm{CIF}_k(\tau) \equiv \Pr(K=k,\ Y \leqslant \tau)$, which measures the absolute risk or incidence of failures from cause k, by time τ. Since any $\mathrm{CIF}_k(\tau)$ depends on *all* the cause-specific hazards (Andersen *et al.*, 2002), the cause-specific hazard for exit k alone is not sufficient to identify the contribution of exit k to the overall mean duration; in extreme cases a covariate, such as treatment status, may have a strong influence on a cause-specific hazard but no effect on the cause-specific CIF (Fine and Gray, 1999; Kyyrä, 2009).

The objective of the current paper is to show how the CIFs can be used to decompose the treatment effect of an intervention, measured in the time domain, into its competing risks components. (This use of $\mathrm{CIF}_k(\tau)$ differs from that proposed in a series of papers by Lo and Wilke (2010) and Lo *et al.* (2016, 2017) who extended the typical cause-specific hazard approach by using estimates of $\mathrm{CIF}_k(\tau)$ together with a copula estimator to determine or bound the unknown marginal survival functions, without imposing independent risks.) We are interested in comparing the difference in the average outcomes of a treated group and a control group at a given threshold, τ weeks after an intervention. Let Y be a random variable representing durations. At time τ, the recorded duration for any individual is either the length of the completed spell, if this is less than τ, or τ if the spell is on going. To carry out the evaluation in this setting I follow Royston and Parmar (2013) and conduct the analysis in terms of restricted mean survival times. The τ-restricted mean lifetime is defined as

$$e(0,\tau) = \int_0^{\tau} S(y)\,\mathrm{d}y$$

where $S(y)=\Pr(Y>y)$. The restricted mean lifetime measures the expected duration up to time τ; $E\{\min(Y,\tau)\}$. To measure the success of an intervention, at time τ, I compare the τ-restricted mean lifetimes of a treatment group T and a control group C. The object of interest is $\Delta(\tau)\equiv e_{\mathrm{T}}(0,\tau)- e_{\mathrm{C}}(0,\tau)$, a treatment effect that is valid under any distribution of time to event and is readily interpretable in the time domain (Royston and Parmar, 2013).

To allow for competing risks, assume that the duration can end in one of K mutually exclusive ways. Andersen (2013) showed that the difference in the τ-restricted mean lifetimes between a treatment group T and a control group C can be written as

$$\Delta(\tau)=\sum_k\left[\int_0^\tau\{\mathrm{CIF}_{k\mathrm{C}}(y)-\mathrm{CIF}_{k\mathrm{T}}(y)\}\mathrm{d}y\right]. \tag{1}$$

The term inside the square brackets provides a measure of the contribution of exit state k to the difference in restricted mean lifetimes. Non-parametric estimation of the CIFs in the presence of right-censored survival data is straightforward by using Kaplan–Meier or Aalen–Johansen methods. $\Delta(\tau)$ is therefore well defined and can be estimated and decomposed, even when the difference in unrestricted means, $\Delta(\infty)$, may be ill determined because of censoring.

The Andersen decomposition, given in equation (1), measures the cause-specific contribution to differences in durations by the differences in the areas under the cause-specific CIFs. (For a comparison of the Andersen approach with other competing risks decompositions see Andersen *et al.* (2013).) However, as noted by Huang *et al.* (2016), analysis based solely on the areas under the CIFs can be confounded by differences in the prevalence rates of the competing events across groups. Using only the difference in areas under the CIFs, it is not possible to distinguish between a programme that resulted in fewer people exiting to work (an incidence effect), but resulted in those who did exit, doing so quicker (a duration effect), and a programme that had no effect on either the duration or incidence of people exiting to work. To account for this, I propose an extension of the Andersen decomposition, which further decomposes the aggregate effect into both a duration effect and an incidence effect.

To derive the extended decomposition, rewrite the difference in restricted means as

$$\begin{aligned}\Delta(\tau)=&\sum_k\left(\mathrm{CIF}_{k\mathrm{T}}(\tau)\left[\tau-\int_0^\tau\frac{\mathrm{CIF}_{k\mathrm{T}}(y)}{\mathrm{CIF}_{k\mathrm{T}}(\tau)}\mathrm{d}y-\left\{\tau-\int_0^\tau\frac{\mathrm{CIF}_{k\mathrm{C}}(y)}{\mathrm{CIF}_{k\mathrm{C}}(\tau)}\mathrm{d}y\right\}\right]\right.\\&\left.+\{\mathrm{CIF}_{k\mathrm{T}}(\tau)-\mathrm{CIF}_{k\mathrm{C}}(\tau)\}\left\{\tau-\int_0^\tau\frac{\mathrm{CIF}_{k\mathrm{C}}(y)}{\mathrm{CIF}_{k\mathrm{C}}(\tau)}\mathrm{d}y\right\}\right)+\underbrace{R_\tau}_{\text{remainder}}\\=&\sum_k\left[\underbrace{\mathrm{CIF}_{k\mathrm{T}}(\tau)\{E(Y_{\mathrm{T}}|K_{\mathrm{T}}=k,Y_{\mathrm{T}}\leqslant\tau)-E(Y_{\mathrm{C}}|K_{\mathrm{C}}=k,Y_{\mathrm{C}}\leqslant\tau)\}}_{\text{duration effect}}\right.\\&\left.+\underbrace{\{\mathrm{CIF}_{k\mathrm{T}}(\tau)-\mathrm{CIF}_{K\mathrm{C}}(\tau)\}E(Y_{\mathrm{C}}|K_{\mathrm{C}}=k,Y_{\mathrm{C}}\leqslant\tau)}_{\text{incidence effect}}\right]+\underbrace{R_\tau}_{\text{remainder}}\end{aligned} \tag{2}$$

where $R_\tau=\tau\{S_{\mathrm{T}}(\tau)-S_{\mathrm{C}}(\tau)\}$. The second equality follows from the fact that

$$E(Y_i|K_i=k,Y_i\leqslant\tau)=\tau-\int_0^\tau\frac{\mathrm{CIF}_{ki}(y)}{\mathrm{CIF}_{ki}(\tau)}\,\mathrm{d}y.$$

(For a proof see Appendix A.)

The first term of the extended decomposition measures the duration effect of state k. If the cause-specific incidences are the same for the control and treatment groups, i.e. $\mathrm{CIF}_{k\mathrm{T}}(\tau)=\mathrm{CIF}_{k\mathrm{C}}(\tau)$, then all of the state k effect will be captured by the duration term, with a negative contribution denoting quicker exits to a given state. The second term measures the incidence effects of each exit state and is given by the differences in the CIFs, weighted by the average duration for group C. If the average duration, conditionally on being in state k, is the same for both groups, but more of the treatment group enter state k, then the first term will be 0 and all of the state k effect will be picked up as a *positive* contribution in the incidence term. How we interpret this component depends on the desirability of the exit state. When exit states are seen as desirable (e.g. work) then a positive incidence effect reflects well on a programme. The opposite is true for exits states that are considered undesirable (e.g. exits into inactivity).

In the above decomposition, $R_\tau=\tau\{S_\mathrm{T}(\tau)-S_\mathrm{C}(\tau)\}$ captures differences in the overall survival rates at time τ. For $\lim_{\tau\to\infty} R_\tau$ to be 0, $S_\mathrm{T}(\tau)-S_\mathrm{C}(\tau)$ must be of order $o(1/\tau)$. Thus an exact decomposition, involving only the competing risks terms, imposes a convergence condition on the survival functions of the treated and control groups. This condition will be automatically satisfied in the absence of censoring. However, it is possible that a policy change may increase or decrease the share of people who never exit, thus possibly leading to differences in higher quantiles, so that R_τ does not vanish. In this case, all three terms will be present in decomposition (2), with the remainder term measuring the effect of on-going spells on the differences in duration.

Further manipulation allows us to rewrite

$$\mathrm{CIF}_{k\mathrm{T}}(\tau)\left[\tau-\int_0^\tau \frac{\mathrm{CIF}_{k\mathrm{T}}(y)}{\mathrm{CIF}_{k\mathrm{T}}(\tau)}\mathrm{d}y-\left\{\tau-\int_0^\tau \frac{\mathrm{CIF}_{k\mathrm{C}}(y)}{\mathrm{CIF}_{k\mathrm{C}}(\tau)}\mathrm{d}y\right\}\right] \\ =\int_0^\tau \{\theta_k \mathrm{CIF}_{k\mathrm{C}}(y)-\mathrm{CIF}_{k\mathrm{T}}(y)\}\mathrm{d}y \qquad (3)$$

where

$$\theta_k=\frac{\mathrm{Pr}_\mathrm{T}(Y\leqslant\tau, D=k)}{\mathrm{Pr}_\mathrm{C}(Y\leqslant\tau, D=k)}=\frac{\mathrm{CIF}_{k\mathrm{T}}(\tau)}{\mathrm{CIF}_{k\mathrm{C}}(\tau)}.$$

Andersen (2013) showed that the total difference in durations due to cause k can be represented as the difference in the areas underneath the cause-specific CIFs. Equation (3) shows, likewise, that the duration effect due to cause k can be represented as the difference in *rescaled* CIFs, where the scaling factor for the control group is simply θ_k. Furthermore, if $\mathrm{CIF}_{k\mathrm{T}}(\tau)=\mathrm{CIF}_{k\mathrm{C}}(\tau)$, decomposition (2) simplifies to that proposed by Anderson (2013); in the absence of any incidence effects both decompositions are identical.

To proceed with the decomposition, it is sufficient to note that well-behaved estimators exist for all the key components of this extended decomposition, $\int_0^\tau \mathrm{CIF}_k(y)\mathrm{d}y$, $\mathrm{CIF}_k(\tau)$ and $S(\tau)$, even in the presence of independent right censoring (Putter *et al.*, 2007). (In practice these components can be easily estimated in Stata by using the `stpci` and `stplost` commands (Overgaard *et al.*, 2015).) Implementing the decomposition involves replacing the unknown $\int_0^\tau \mathrm{CIF}_k(y)\mathrm{d}y$, $\mathrm{CIF}_k(\tau)$ and $S(\tau)$ with their corresponding estimators. By varying τ, the decomposition enables us to examine how the contribution of the alternative states varies over the duration of a spell.

3. Application: changes to unemployment benefit

I use the above decomposition to evaluate the effect of a substantial cut to Jobseeker's Allowance (JA) in Ireland in 2009. JA is a means-tested benefit, which is paid indefinitely to those who are unemployed, provided that the claimant meets the conditions of eligibility. In response to the fiscal crisis that occurred in Ireland with the onset of the Great Recession, the Irish Government introduced a substantial cut to JA payments for young claimants. Before 2009, all claimants were entitled to €204 per week. On April 29th, 2009, claimants aged 18 years had their weekly rate cut to €100. The benefit cut applied only to new claimants; those who started their claim before April 29th continued to receive the higher rate. As a result, a comparison of those who entered unemployment just before and just after the legislation can be used to identify the treatment effect. (Doris *et al.* (2017) found no evidence of an anticipation effect before the introduction of the legislation.)

Doris *et al.* (2017) analysed this benefit cut in detail, using a regression discontinuity design, and found that claimants subject to the cut had unemployment durations that were 50% shorter than the durations of those in the control group. Given that the stated motivation for these cuts was to 'ensure that young people are better off in education, employment or training than claiming' (`http://www.welfare.ie/en/pressoffice/Pages/pr231013.aspx`), it seems appropriate to consider, in detail, the relative importance of alternative exit states in explaining the aggregate effect.

4. Data

To carry out the analysis, I use the longitudinal jobseekers' database that is provided by the Department of Social Protection. This is an administrative data set, covering every claimant who has received a jobseeker's or one-parent family payment since 2007. The data provide administrative records for the start and end date of every new claim, allowing me to establish both the exact start date and the duration of claiming unemployment benefits for the entire population of new JA claims initiated between 2007 and 2014. I denote claim duration as unemployment duration in what follows, although the former is only a subset of the latter. Throughout the analysis, I measure unemployment duration in weeks. The administrative data also record the exit state into which the claimant entered, after the unemployment spell ended. When considering the competing risks model, I examine five different exit states: training, work, education, inactivity and 'other'. (The 'other' category is a residual grouping comprising, primarily, durations for which the caseworkers could not establish an exit state for the spell.) The data on exit states come from information that was provided by caseworkers, following the closure of a claim. Information on training and education schemes are directly available from administrative records. Some of the information on employment is likely to have been provided by the claimant. I conducted several cross-checks to verify the accuracy of these data. For instance, in cases where it was reported that the claimant exited to employment I checked official administrative records, to see whether there was an employment tax record for the claimant in the year that they exited unemployment. This was so for the majority of claimants. In a small number of cases, some of those in the 'other' category had tax earnings records that suggested that they had exited to work. I experimented by allocating these claimants to the work category, but this had little effect on the results reported.

To identify the treatment effect, I compare the durations of those beginning an unemployment spell in April 2009 (who received the higher benefit payment) with those beginning a spell in May 2009 (and who received the substantially reduced payment). Table 1 provides summary statistics, separately for those in the control and treatment groups. The first row of Table 1 reports the average unemployment duration, by month of entry and gender. For all groups, the average unemployment duration exceeded 1 year, highlighting the very poor labour market prospects facing young people in Ireland during the Great Recession. However, those who entered after the

Table 1. Summary statistics by month of entry to unemployment

	Results for males		*Results for females*	
	May	*April*	*May*	*April*
Average unemployment duration (weeks)	66.53	99.98	61.20	92.00
Proportion exiting to alternative labour market states				
Proportion exiting to training	0.35	0.31	0.31	0.30
Proportion exiting to work	0.31	0.36	0.36	0.21
Proportion exiting to education	0.05	0.06	0.06	0.11
Proportion exiting to other	0.26	0.25	0.20	0.19
Proportion exiting to inactive	0.03	0.03	0.06	0.18
N	428	377	297	210

benefit cut had unemployment durations that are approximately 30 weeks shorter, on average, than those entering earlier.

Of particular interest in my study is the labour market states into which claimants exited. The next five rows of Table 1 provide the proportions exiting to each of the alternative states. For men, we see that three exit states dominated, namely training, work and other. Furthermore, the proportion exiting to each state is similar for those entering in April and in May. This suggests that the aggregate duration effect, for men, is driven by differences in duration, conditional on exit states. The fourth column of Table 1 shows that, for women subject to the benefit cut, the proportion exiting to each state is similar to those of men and is dominated by the same three exit states. However, for women in receipt of the higher benefit, the pattern is notably different. The difference, relative to those subjected to the lower benefit regime, is driven by a bigger proportion of the treatment group exiting to work, and fewer exiting to inactivity and education. Although differences in state-specific durations appear to be driving the male effect, the results in the last two columns suggest that differences in incidences play a larger role for women. In the remainder of the paper, I use the decomposition that was outlined in Section 3 to examine these issues formally.

5. Results

Since all of my sample had completed their unemployment spell 7 years after the intervention, the remainder term in equation (2) is 0 by the end of my sample period. I therefore begin with an exact decomposition of the overall means. The results are given in Table 2. (The standard errors that are reported for the decomposition terms are estimated by using a bootstrap procedure with 1000 replications.) The first row reports the overall treatment effects. In keeping with Doris *et al.* (2017), I estimate a large and statistically significant duration response to the benefit cut for both men and women. The size of the aggregate effect is similar across genders; males subject to the benefit cut have unemployment durations that are over 33 weeks shorter than those in receipt of the higher benefit, whereas the effect for females is approximately 31 weeks. (In practice some of the treatment group were exempted from benefit cut on family or health grounds. I have also estimated the model by using month of treatment as an instrument to control for this. Although the overall treatment effect increases somewhat there is still little evidence of a significantly different aggregate effect for men and women.)

Table 2. Competing risks decomposition†

Decomposition	*Results (weeks) for males*			*Results (weeks) for females*		
	Overall decomposition (1)	*Duration effect (2)*	*Incidence effect (3)*	*Overall decomposition(4)*	*Duration effect (5)*	*Incidence effect (6)*
Treatment effect		−33.50 [−44.78 : −22.00]			−30.79 [−45.54 : −16.67]	
Training	−6.35	−10.20	3.85	−3.04	−3.63	0.59
	[−14.36 : 2.89]	[−18.14:−2.86]	[−1.98 : 9.97]	[−11.91 : 5.55]	[−10.40:4.03]	[−5.56:6.17]
Work	−14.80	−10.15	−4.65	−1.42	−16.60	15.18
	[−23.72:−5.67]	[−17.03:−3.35]	[−11.42:2.44]	[−11.82:8.08]	[−28.79:−5.32]	[6.46:25.31]
Education	−2.83	−1.94	−0.90	−5.24	−1.00	−4.25
	[−7.19:1.87]	[−4.64:0.42]	[−5.08:3.18]	[−11.22:0.04]	[−3.74:1.53]	[−9.37:0.36]
Other	−6.94	−8.06	1.12	−7.10	−7.82	0.72
	[−14.48:0.34]	[−14.05:−2.48]	[−4.12:6.26]	[−15.16:0.27]	[−14.82:−1.63]	[−5.30:7.00]
Inactive	−2.57	−1.84	−0.73	−13.99	0.04	−14.03
	[−7.53:2.08]	[−4.38:0.48]	[−5.79:3.77]	[−22.93:−5.16]	[−2.89:3.04]	[−21.76:−6.43]

†Non-parametric 95% bootstrap confidence intervals are in parentheses (number of bootstrap replications, 1000).

Columns (1) and (4) of Table 2 present the total contribution of each exit state to the overall treatment effect, making no distinction between incidence and duration effects. Looking at the results for males, we see that no single exit state dominates. Whereas the contribution of exits to education and inactivity are small, the other three exit states all contribute substantially to the overall effect. The shorter overall duration for males is therefore a result of exits to training, work and other. Looking at the results for females in column (4), we see that exits to work contribute the least to the overall female effect. In contrast, exits to inactivity, although not important for men, are the dominant channel for women. The fact that very similar aggregate effects for men and women operate through very different channels highlights the potential value of the competing risks decomposition that is developed in this paper.

The contributions of the duration and incidence effects, to these overall state-specific components, are given in columns (2) and (3) of Table 2 for men, and columns (5) and (6) for women. Looking at the duration effect for men, we see that, conditionally on exiting to a given state, those in receipt of the lower benefit exited quicker than those in the control group. For instance, the results show that, if the overall proportion exiting to training had been identical, the average durations of those in the treatment group would still have been approximately 10 weeks shorter, as a result of the quicker exits to training. The duration effect is particularly evident for the three main effects that were identified earlier: training, work and other. The incidence effects, in contrast, show that differences in the likelihood of exiting to a given state are relatively less important for men.

The duration effects for women in Table 2 are quite similar to those reported for men. Conditionally on exiting to a given state, those in receipt of the lower benefit exit quicker than those in the control group. This is particularly true for exits to work, where the duration effect is larger for women than for men. This is despite the fact that exits to work did not appear to be important for women in the overall decomposition. This is because, in addition to the duration effect, the work incidence effect is *positive* and large for women, implying that women who are subject to the benefit cut are substantially *more* likely to enter work. Both of these work effects would be considered desirable but offset each other in the aggregate decomposition.

The importance of distinguishing between duration and incidence effects for women is also evident when we consider exits to inactivity. The estimate in the final row of column (5) shows little difference in the timing of exits to inactivity, for women in the treatment and control groups. The large overall effect of inactivity that is apparent in column (4) is therefore driven entirely by the incidence effect; women subject to the benefit cut are much *less* likely to exit to inactivity.

Exits to inactivity for women can also be used to illustrate graphically the value of the new decomposition. Fig. 1(a) compares the inactivity CIFs for women in the control and treatment groups. A comparison of the areas under these curves corresponds to the decomposition that was proposed by Andersen (2013). Such a comparison clearly reveals substantial differences between the two groups. However, although the differences in incidences can be inferred from this graph, it is impossible, from this comparison alone, to determine the relative importance of duration and incidence effects. Fig. 1(b) provides the graph for the *conditional duration* effect, using the rescaling adjustment that is suggested in equation (3). The similarity of these two curves immediately implies that the overall effect that is observed in Fig. 1(a) is driven entirely by the incidence component.

The decomposition that was presented in Section 3 also enables us to examine the evolution of the overall treatment effect, by varying the follow-up threshold τ. The results in Tables 3–5 provide the decomposition at three thresholds: 2 years, 4 years and 6 years after the intervention. Looking at Table 3 we see that, for both men and women, the restricted mean treatment effect

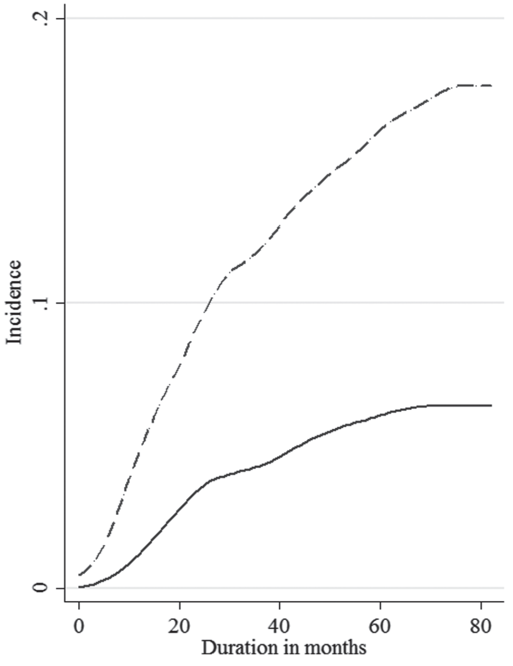

(a)

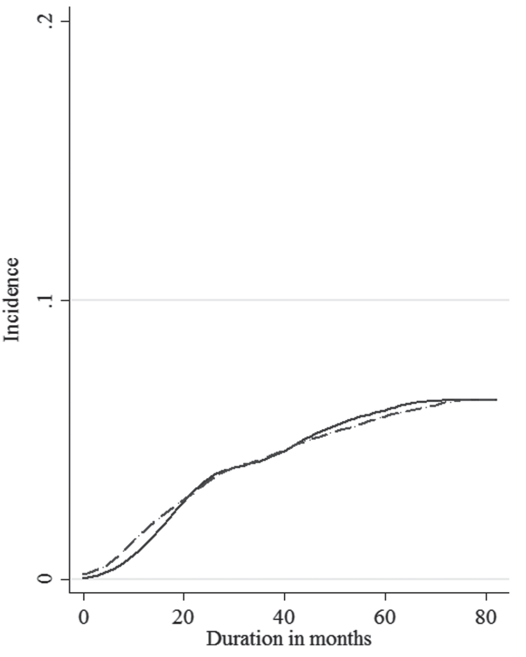

(b)

Fig. 1. Graphical representation of (a) total effect and (b) duration effect accounted for by exits to inactivity for females: ———, May; — —, April

at 2 years is large and statistically significantly different from 0. Furthermore, in keeping with the overall final effect, there is little difference in the magnitude of the 2-year treatment effect between men and women. This is also true of the estimated treatment effects at the 4-year and 6-year thresholds.

The results in Table 3 also show that the incidence effects for women, which were highlighted in Table 2, are already evident at the 2-year threshold; even in the early stages of the unemployment spell women who were subjected to the benefit cut are more likely to exit to work, and less likely to exit to inactivity. For men the dynamics are different. In Table 2 we saw that the overall effect for men was driven primarily by the training and work duration effects. However, from Tables 3–5, we see that these patterns take some time to emerge; the significant training duration effect is only evident at the 4-year threshold, whereas the work duration effect is not evident until the 6-year threshold. In the early stages of the spell, it appears that males who were subjected to the benefit cut reacted mostly by increasing participation in training schemes. It follows from this that the results of any evaluation, such as this, must take into account the timing of the study. In my example, a short-run follow-up would reveal only a training incidence effect, while masking the work duration effect that is evident in the longer run.

6. Conclusion

In this paper, I consider how to deal with competing risks when the outcome of interest is the difference in average duration between two groups. I propose a competing risks decomposition, which identifies the contribution of each exit state to the overall difference in duration between

Table 3. Competing risks decomposition of restricted means at 2 years†

Decomposition	*Results (weeks) for males*			*Results (weeks) for females*		
	Overall decomposition (1)	*Duration effect (2)*	*Incidence effect (3)*	*Overall decomposition (4)*	*Duration effect (5)*	*Incidence effect (6)*
Restricted mean difference (2 years)		−13.02 [−18.00:−8.12]			−13.09 [−19.41:−6.99]	
Training	3.07	−0.15	3.22	0.61	−0.28	0.87
	[0.47:5.76]	[−2.20:1.78]	[1.08:5.34]	[−2.63:3.60]	[−2.28:1.76]	[−1.73:3.24]
Work	1.04	0.36	0.67	5.18	−1.07	6.20
	[−1.55:3.74]	[−1.42:2.28]	[−1.09:2.60]	[1.82:8.44]	[−4.73:2.68]	[3.41:9.53]
Education	0.25	−0.20	0.45	−1.25	−0.06	−1.19
	[−1.17:1.62]	[−0.99:0.53]	[−0.84:1.98]	[−3.51:0.72]	[−0.94:0.84]	[−3.05:0.41]
Other	1.16	0.75	1.91	0.69	−0.77	1.43
	[−1.26:3.53]	[−2.38:0.80]	[−0.30:4.00]	[−2.13:3.40]	[−3.31:1.40]	[−0.74:4.24]
Inactive	0.02	−0.22	0.23	−2.27	0.47	−2.74
	[−0.57:0.53]	[−4.99:6.17]	[−5.66:5.34]	[−5.09:0.35]	[−0.14:1.24]	[−5.30:−0.53]
Remainder	−18.55			−16.05		
	[−24.86:−12.09]			[−23.99:−7.59]		

†Non-parametric 95% bootstrap confidence intervals are in parentheses (number of bootstrap replications, 1000).

Table 4. Competing risks decomposition of restricted means at 4 years†

Decomposition	*Results (weeks) for males*			*Results (weeks) for females*		
	Overall decomposition (1)	*Duration effect (2)*	*Incidence effect (3)*	*Overall decomposition (4)*	*Duration effect (5)*	*Incidence effect (6)*
Restricted mean difference (4 years)		−26.54 [−35.68:−17.43]			−25.97 [−37.35:−14.42]	
Training	−2.05 [−6.97:2.42]	−5.57 [−9.72:−1.69]	3.53 [−0.22:7.68]	−2.67 [−9.04:3.22]	−3.34 [−8.31:1.58]	0.67 [−4.03:5.06]
Work	−2.92 [−8.99:2.83]	−2.65 [−6.94:1.19]	−0.26 [−4.65:3.92]	4.99 [−0.22:10.63]	−4.56 [−10.80:1.71]	9.55 [5.16:15.20]
Education	−1.16 [−4.01:2.00]	−1.04 [−2.74:0.56]	−0.12 [−2.77:2.91]	−2.56 [−7.03:1.38]	0.15 [−1.84:2.17]	−2.71 [−6.58:0.28]
Other	−0.04 [−4.36:4.38]	−2.23 [−5.53:1.06]	2.19 [−1.19:5.67]	−2.29 [−7.74:2.89]	−3.83 [−8.10:0.24]	1.54 [−2.95:6.07]
Inactive	0.01 [−2.19:0.51]	−0.64 [−2.12:0.51]	0.64 [−1.41:2.80]	−8.31 [−13.97:−1.98]	0.16 [−1.34:1.81]	−8.47 [−13.66:−3.17]
Remainder	−20.38 [−30.38:−11.14]			−15.13 [−26.49:−3.82]		

†Non-parametric 95% bootstrap confidence intervals are in parentheses (number of bootstrap replications, 1000).

Table 5. Competing risks decomposition of restricted means at 6 years†

Decomposition	*Results (weeks) for males*			*Results (weeks) for females*		
	Overall decomposition (1)	*Duration effect (2)*	*Incidence effect (3)*	*Overall decomposition (4)*	*Duration effect (5)*	*Incidence effect (6)*
Restricted mean difference (6 years)		−33.06 [−44.11:−21.63]			−30.67 [−45.08:−16.53]	
Training	−6.04	−9.70	3.66	−3.05	−3.63	0.58
	[−13.18:0.99]	[−16.55:−3.37]	[−1.58:9.32]	[−11.91:5.55]	[−10.40:4.03]	[−5.56:6.17]
Work	−10.26	−7.36	−2.95	−0.71	−14.67	13.95
	[−19.17:−2.21]	[−14.00:−1.09]	[−9.29:3.36]	[−10.07:7.71]	[−25.67:−4.98]	[6.24:22.83]
Education	−1.95	−1.42	−0.53	−3.65	−0.30	−3.35
	[−5.86:2.32]	[−3.81:0.91]	[−4.19:3.26]	[−8.55:0.81]	[−2.38:1.92]	[−7.74:0.05]
Other	−6.82	−7.92	−0.81	−3.95	−5.29	1.33
	[−13.67:−2.64]	[−13.27:−2.65]	[−6.54:4.92]	[−10.57:2.71]	[−11.19:0.18]	[−3.76:6.72]
Inactive	−1.61	−1.51	−0.09	−13.99	0.04	−14.03
	[−5.76:2.08]	[−3.79:0.49]	[−4.21:3.91]	[−22.93:−5.16]	[−2.89:3.04]	[−21.76:−6.43]
Remainder	−6.38			−5.33		
	[−13.22:0.46]			[−12.32:1.23]		

†Non-parametric 95% bootstrap confidence intervals are in parentheses (number of bootstrap replications, 1000).

the two groups, distinguishing between differences in duration, conditionally on an exit state, and differences in the absolute risk of exits to each state.

I use the decomposition to examine the effect of a 50% cut to unemployment benefit in Ireland. The aggregate effect of the benefit cut is substantial for both men and women, with unemployment duration falling by 33 weeks for men and 31 weeks for women. However, the decomposition reveals substantial gender differences in the channels through which this effect operates, with the incidence effect substantially more important for women than for men. These differences may also have implications for the long-run effects of the benefit cut, which are not examined in this paper. The increase in the number of young women exiting work after the benefit cut may strengthen their life-long attachment to the labour market, resulting in additional benefits, which are not apparent when we focus only on the duration of unemployment at the time of the cut.

Acknowledgements

I am grateful to Terry Corcoran (Department of Social Protection) for providing the Department of Social Protection longitudinal data that were used in this analysis and for many useful discussions in relation to these data. I am also grateful to Aedín Doris and Olive Sweetman for key discussions and insights on earlier versions of this paper.

Appendix A

Theorem 1. Let Y be any random variable:

$$E(Y|K=k,\ Y\leqslant\tau)=\tau-\int_0^{\tau}\frac{\mathrm{CIF}_k(y)}{\mathrm{CIF}_k(\tau)}\mathrm{d}y$$

Proof. For any random variable X

$$E(X)=\int_0^{\infty}S(x)\mathrm{d}x \qquad S(x)=\Pr(X>x).$$

Specifically,

$$E(Y|K=k,\ Y\leqslant\tau)=\int_0^{\infty}\Pr(Y\geqslant u|K=k,\ Y\leqslant\tau)\,\mathrm{d}u.$$

Using the conditional probability law we can rewrite this as

$$\int_0^{\infty}\frac{\Pr(Y\geqslant u,\ Y\leqslant\tau,\ K=k)}{\Pr(Y\leqslant\tau,\ K=k)}\,\mathrm{d}u.$$

By definition of $\mathrm{CIF}_k(\tau)$ this can be written as

$$\frac{1}{\mathrm{CIF}_k(\tau)}\int_0^{\infty}\Pr(Y\geqslant u,\ Y\leqslant\tau,\ K=k)\,\mathrm{d}u$$
$$=\frac{1}{\mathrm{CIF}_k(\tau)}\left\{\int_0^{\tau}\Pr(Y\geqslant u,\ Y\leqslant\tau,\ K=k)\,\mathrm{d}u+\int_{\tau}^{\infty}\Pr(Y\geqslant u,\ Y\leqslant\tau,\ K=k)\,\mathrm{d}u\right\}.$$

However, since the last term above is 0 we can rewrite this as

$$=\frac{1}{\mathrm{CIF}_k(\tau)}\int_0^{\tau}\Pr(Y\geqslant u,\ Y\leqslant\tau,\ K=k)\,\mathrm{d}u.$$

Using the fact that $\Pr(Y>y,\ K=k)=1-\Pr(Y<y,\ K=k)-\Pr(K\neq k)=\Pr(K=k)-\Pr(Y<y,\ K=k)$ this equals

$$
\begin{aligned}
&\frac{1}{\mathrm{CIF}_k(\tau)}\int_0^{\tau}\{S_k(u)-S_k(\tau)\}\,\mathrm{d}u \qquad\qquad S_k(y)\equiv \Pr(Y>y,\ K=k)\\
&\quad=\frac{1}{\mathrm{CIF}_k(\tau)}\left\{\int_0^{\tau}\Pr(K=k)\,\mathrm{d}u-\int_0^{\tau}\mathrm{CIF}_k(u)\,\mathrm{d}u-\int_0^{\tau}S_k(\tau)\,\mathrm{d}u\right\}\\
&\quad=\frac{1}{\mathrm{CIF}_k(\tau)}\left\{\tau\Pr(K=k)-\int_0^{\tau}\mathrm{CIF}_k(u)\,\mathrm{d}u-\tau\,S_k(\tau)\right\}.
\end{aligned}
$$

Again using the fact that $\Pr(Y>\tau,\ K=k)=\Pr(K=k)-\Pr(Y<\tau,\ K=k)$ this equals

$$
\begin{aligned}
&\frac{1}{\mathrm{CIF}_k(\tau)}\left(\tau\Pr(K=k)-\int_0^{\tau}\mathrm{CIF}_k(u)\,\mathrm{d}u-[\tau\{\Pr(K=k)-\mathrm{CIF}_k(\tau)\}]\right)\\
&\quad=\frac{1}{\mathrm{CIF}_k(\tau)}\left\{\tau\,\mathrm{CIF}_k(\tau)-\int_0^{\tau}\mathrm{CIF}_k(u)\,\mathrm{d}u\right\}\\
&\quad=\tau-\int_0^{\tau}\frac{\mathrm{CIF}_k(u)}{\mathrm{CIF}_k(\tau)}\,\mathrm{d}u
\end{aligned}
$$

as required.

Note that if $K=1$ so that there is only one exit state this simplifies to the standard expression for the conditional mean:

$$E(Y\,|\,Y\leqslant\tau)=\tau-\int_0^{\tau}\frac{\mathrm{CIF}(y)}{\mathrm{CIF}(\tau)}\mathrm{d}y=\int_0^{\tau}\frac{S(y)-S(\tau)}{1-S(\tau)}\mathrm{d}y.$$

References

Andersen, P. K. (2013) Decomposition of number of life years lost according to causes of death. *Statist. Med.*, **32**, 5278–5285.

Andersen, P. K., Abildstrom, S. and Rosthøj, S. (2002) Competing risks as a multi-state model. *Statist. Meth. Med. Res.*, **11**, 203–215.

Andersen, P. K., Canudas-Romo, V. and Keiding, N. (2013) Cause specific measures of life years lost. *Demog. Res.*, **29**, 1127–1152.

Andersen, P. K., Geskus, R. B., De Witte, T. and Putter, H. (2012) competing risks in epidemiology: possibilities and pitfalls. *Int. J. Epidem.*, **41**, 861–870.

Austin, P. C. and Fine, J. P. (2017) Accounting for competing risks in randomized controlled trials: a review and recommendations for improvement. *Statist. Med.*, **36**, 1203–1209.

van den Berg, G. J., van der Klaauw, B. and Van Ours, B. (2014) Punitive sanctions and the transition rate from welfare to work. *J. Lab. Econ.*, **22**, 211–241.

Dolton, P. and O'Neill, D. (1996) Unemployment duration and the restart effect: some experimental evidence. *Econ. J.*, **106**, 387–400.

Doris, A., O'Neill, D. and Sweetman, O. (2017) Does reducing unemployment benefits during a recession reduce youth unemployment?: Evidence from a 50% cut in unemployment assistance. *Discussion Paper 10727*. IZA Institute of Labor Economics, Bonn.

Fine, J. P. and Gray, R. J. (1999) A proportional hazards model of the subdistribution of a competing risk. *J. Am. Statist. Ass.*, **94**, 496–509.

Han, A. and Hausman, J. A. (1990) Flexible parametric estimation of duration and competing risk models. *J. Appl. Econmetr.*, **5**, 1–28.

Huang, C.-Y, Wang, C. and Wang, M.-C. (2016) Nonparametric analysis of bivariate gap time with competing risks. *Biometrics*, **72**, 780–790.

Katz, L. F. and Meyer, B. D. (1990) The impact of the potential duration of unemployment benefits on the duration of unemployment. *J. Publ. Econ.*, **41**, 45–72.

Kyyrä, T. (2009) Marginal effects for competing risks models with piecewise constant hazards. *Oxf. Bull. Econ. Statist.*, **71**, 539–565.

Lo, S., Stephan, G. and Wilke, R. (2016) Estimating the latent effect of unemployment benefits on unemployment duration. *Discussion Paper 6650.* IZA Institute of Labor Economics, Bonn.

Lo, S., Stephan, G. and Wilke, R. (2017) Competing risks copula models for unemployment duration: an application to a German Hartz reform. *J. Econmetr. Meth.*, **6**, 1–20.

Lo, S. M. S. and Wilke, R. A. (2010) A copula model for dependent competing risks. *Appl. Statist.*, **59**, 359–376.
McVicar, D. (2008) Job search monitoring intensity, unemployment exit and job entry. *Lab. Econ.*, **15**, 1451–1468.
Overgaard, M., Andersen, P. K. and Parner, E. T. (2015) Regression analysis of censored data using pseudo-observations: an update. *Stata J.*, **15**, 809–821.
Putter, H., Fiocco, M. and Geskus, R. B. (2007) Tutorial in biostatistics: competing risks and multi-state models. *Statist. Med.*, **26**, 2389–2430.
Royston, P. and Parmar, M. (2013) Restricted mean survival time: an alternative to the hazard ratio for the design and analysis of randomised trials with a time-to-event outcome. *BMC Med. Res. Methodol.*, **13**, article 152.