

European Union Politics

<http://eup.sagepub.com>

Temporal Change and the Process of European Union Decision-Making

Christopher Zorn

European Union Politics 2007; 8; 567

DOI: 10.1177/1465116507082815

The online version of this article can be found at:

<http://eup.sagepub.com>

Published by:



<http://www.sagepublications.com>

Additional services and information for *European Union Politics* can be found at:

Email Alerts: <http://eup.sagepub.com/cgi/alerts>

Subscriptions: <http://eup.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.co.uk/journalsPermissions.nav>

Citations <http://eup.sagepub.com/cgi/content/refs/8/4/567>



European Union Politics

DOI: 10.1177/1465116507082815

Volume 8 (4): 567–576

Copyright© 2007

SAGE Publications

Los Angeles, London, New Delhi
and Singapore

Temporal Change and the Process of European Union Decision-Making



Christopher Zorn

Pennsylvania State University, USA

Introduction

Survival analysis – also variously referred to as ‘event history analysis’, ‘duration analysis’ and ‘reliability analysis’ – is an increasingly widespread approach to the quantitative study of political phenomena. Models for survival data can be traced to early work in biostatistics, and before that to life tables developed in the actuarial sciences; the latter are traceable at least to 18th-century work by Bernoulli, D’Alembert and their contemporaries.¹ Yet it was not until the 20th century that statisticians, particularly those interested in engineering, operations research and medicine, began developing methods for analyzing time-to-event data. The development of such models took a two-pronged path, with scholars in engineering and other physical sciences gravitating toward (and thus developing) parametric models, while those in epidemiology and medicine focused to a far greater extent on Cox’s (1972) semi-parametric approach.²

Since the initial application of survival models in the 1970s, their use in the social and behavioral sciences has grown tremendously. That growth can be traced to a number of related causes. Certainly increases in computing power and the widespread availability of software to estimate such models have played a role, as have the greater availability, reliability and extent of time-to-event data in the social sciences. Equally important, however, has been a renewed substantive focus on social and political *processes*, and in particular on how preferences and institutions interact to shape decision-making.

The articles on EU decision-making in this and the forthcoming issue of *European Union Politics* reflect that renewed focus, taking as their phenomenon of interest the speed with which the EU passes legislation. The

authors offer several improvements – theoretical, methodological and empirical – on existing studies, many of which revolve around the importance of the ‘proportional hazards’ assumption common to such models. Here, I shall address some aspects of those improvements and discuss how analysts might integrate these insights with theories to develop and test more complete models of EU decision-making.

Time-varying covariates

Whereas early studies of event histories in political science tended to adopt parametric approaches, more recent work has increasingly moved toward Cox’s (1972) proportional hazards model as the model of choice. As noted above, that model was developed in the context of biostatistics and life-table analysis in epidemiology. In such fields, the paradigmatic study is the double-blind clinical trial: subjects with some medical condition are randomly selected into treatment and control (placebo) groups, the treatment is administered, and the subjects are then followed and the time to the event of interest (e.g. recurrence of the condition) is recorded. Such trials contemplate a relatively simple model, where the hazard of the event of interest is a function of a single, binary covariate:

$$h_i(t) = f\{\beta_0 + \beta_1 X_i\}. \quad (1)$$

In this model, X denotes the treatment indicator, coded 1 for treated individuals and 0 for those in the placebo group; the usual expectation is that $\beta_1 < 0$ (i.e. that the treatment reduces the hazard of recurrence). Note that, as written here, the model in equation (1) is a general one, encompassing both parametric and semi-parametric alternatives. Nonetheless, equation (1) implies (among other things) that the value of the covariate X for a given subject i is fixed in the study, i.e. either subjects receive the treatment (placebo) only once, or those selected to receive the treatment (placebo) continue to receive it over the life of the study.

The simplest extension of this model is the two-period crossover design (Jones and Kenward, 2003: chapter 2). In this design, the follow-up is divided into two periods: period 1 corresponds to the design above, whereas in period 2 the sample of subjects initially receiving the treatment is ‘switched’ to the placebo, and vice versa.³ Crossover designs thus represent the most basic form of ‘time-varying covariate’: each subject now has both $X = 0$ and $X = 1$, with the value of X varying across the two periods. We can re-express equation (1) to capture this variation by writing:

$$h_i(t) = f\{\beta_0 + \beta_1 X_{it}\}, \quad (2)$$

where the new subscript t on X denotes that the value of the covariate now depends both on the identity of the subject and on the time of measurement.

Importantly, the specification in (2) retains a key assumption of (1): that the effect of the treatment on the hazard is invariant across the two periods. Put differently, (2) assumes that the period in which the treatment is administered is irrelevant to its impact on the hazard of event recurrence. If this is not the case – if, for example, the treatment is more effective for those receiving it in period 1 than for those receiving it in period 2 – the result is what is often termed a *treatment* \times *period* interaction (e.g. Armitage and Hills, 1982). Although notations vary, we can write such an interaction in its most general form as:

$$h_i(t) = f\{\beta_0 + \beta_1 X_{it} + \beta_2 [X_{it} \times g(t)]\}. \quad (3)$$

Here, the extent of X 's effect on the hazard of recurrence is allowed to change over time. For $g(t)$ equal to a step function (say, $g(t) = 0$ in period 1 and $g(t) = 1$ in period 2), the model in (3) captures the possibility of treatment & period interactions; a test for $\beta_2 = 0$ corresponds to a test for such interactive effects.⁴

The two-period crossover design is a special (and relatively simple) example of a covariate that changes over time. Note, however, that the mere fact of that change does not, by itself, necessarily imply that the variable's effects on the hazard will also vary over time; it is entirely possible that no treatment \times period interaction is present, and that the effect of the drug on the hazard of recurrence is identical, irrespective of when it is administered. Nor is it the case that such temporal variation is necessary for a covariate's effects to change over time; one can easily imagine a case where the influence of X on the hazard declined (or rose) over the course of the study's follow-up period, even in the absence of a crossover design. This fact will become important in the discussion of non-proportionality, below.

Baseline hazards and (non-)proportionality

As the authors of the present studies are well aware, a key assumption of most widely used survival models – including Cox's model and the exponential and Weibull parametric models – is that of the proportionality of covariate effects. In the language of my example above, proportionality requires that the effect of a covariate on the hazard be proportional; formally, this means that

$$\frac{h(t | X = 0)}{h(t | X = 1)} = c \forall t. \quad (4)$$

Proportionality is 'built in' to those models as a function of how the covariates impact the hazard (that is, exponentially). As a practical matter, the marginal effect of a covariate (which can be thought of as how 'far apart' $h(t | X = 0)$ and $h(t | X = 1)$ are) will necessarily vary over time – even if the covariate's effect is proportional – so long as the shape of the baseline hazard is anything other than constant over time.

Non-proportionality of covariate effects, then, amounts to any deviation from equation (4), and can arise owing to several phenomena. The simplest case is when the effect of the covariate on the underlying hazard of the event changes over time, e.g. in the manner described in equation (3). This would be the case, for example, if patients developed resistance to a drug's effects – such that its curative power declined the longer the patient took the drug – or, conversely, if the drug accumulated in the patient's body, leading to larger effects over time. In the former instance, the hazards for the treatment and placebo groups would be converging over time, whereas in the latter they would diverge.⁵

A second, related (but somewhat more subtle) cause of non-proportionality occurs when the effect of a covariate is to change the shape of the baseline hazard. Consider, for example, a drug that, rather than lowering the hazard of disease recurrence, instead caused the shape of the baseline hazard of the event to change from being 'flat' (that is, constant over time) to declining over time. This is presented in Figure 1, where the solid line represents the 'baseline' hazard (i.e. that for $X = 0$, the placebo group) and the dashed lines illustrate two alternative covariate influences. The common way of thinking about a covariate's effect – illustrated by the short-dashed line – is as causing a discrete change in the hazard. Here, the effect of the treatment ($X = 1$) is to lower the hazard of recurrence; in Figure 1, that effect is illustrated as being proportional, in that the ratio $h(t | X = 0)/h(t | X = 1)$ remains constant over time. In contrast, the long-dashed line illustrates an alternative effect: the mean/average hazard for the treated group is the same as for the control/baseline group, but the treatment causes the hazard to go from 'flat' to decreasing over time. Note that, in the latter case, the effect of the treatment on the (average) *level* of the hazard is zero; nonetheless, its effect on the shape of the hazard causes the condition in equation (4) to be violated.

The issue of deviations from proportionality is a well-researched one in biostatistics, and the development and widespread availability of numerous tests for violations of proportionality mean that applied researchers have at their disposal a range of tools for identifying such violations. But, although detecting the *presence* of non-proportionality is relatively simple, as a practical matter, ascertaining the *source* of non-proportionality can be more




Figure 1 Potential covariate effects: Constant hazards.

Notes: The solid line is the 'baseline' hazard ($X = 0$). The short-dashed line is the hazard for the treatment group ($X = 1$) when the effect of X is on the level of the hazard. The long-dashed line is the hazard for the treatment group ($X = 1$) when the effect of X is to change the 'shape' (but not the level) of the baseline hazard.

difficult. In particular, covariates that have an effect on the shape of the hazard often appear very similar to those that (non-proportionally) affect its overall level.

As an illustration of this similarity, Figure 2 plots three hypothetical pairs of hazards⁶ defined by their relationship to a binary covariate X . The first, panel (a), illustrates the case where the effect of the covariate is proportional; the baseline (placebo group) hazard (the smooth line, denoting $X = 0$) is decreasing and the hazard for the treated group (the dashed line, for $X = 1$) is also decreasing, such that the ratio of the two remains constant over time. Panel (b) plots the relative hazards for treatment and placebo groups in the more familiar non-proportional case, when the effect of the treatment increases (log-)linearly over time. Note that the influence of X increases, from negligible at the beginning of the follow-up period to substantial at its end. Finally, panel (c) illustrates the apparent effect of a covariate that (here, positively) influences the shape of the hazard but has no influence on its level. Panels (b) and (c) are very similar; in both, there is clear evidence that X 's




Figure 2 Three covariate effects on $h(t)$.

Notes: The figures plot hypothetical hazards for $X = 0$ (smooth line) and $X = 1$ (dashed line). In panel (a), the hazards are proportional. In panel (b), the (negative) effect of X is nonproportional (increasing over time). In panel (c), X has no effect on the level of the hazard itself, but alters the shape of the baseline hazard.

influence is non-proportional, but the root cause of that non-proportionality is different for each.

Thinking about covariates' effects on the shape of the underlying hazard of interest requires that we reconceptualize the baseline hazard as something other than a 'nuisance', and instead consider its potential substantive meaning and/or significance. In particular, it means that we need to consider theorizing about the potential determinants of the shape of the baseline hazard and how covariates might alter that shape. To do so, in turn, suggests a need to theorize not simply about what makes events likely to happen sooner or later but also about how the process giving rise to those events might change under different exogenous conditions.

Implications for EU decision-making

Contemplating the shape of the baseline hazard forces analysts to think not just about what factors make events more or less likely to occur but also about

how and why that likelihood changes during the process under study. In the context of EU decision-making, it becomes necessary to ask, as the decision-making process goes on, if – and, more importantly, why – the conditional odds of passage increase, decrease or remain the same.

Consider but one example: the influence of membership size and veto players on decision-making speed. Building on work by Tsebelis (e.g. 2002) and others, the authors in this and the forthcoming issue draw a connection between (for example) the level of disagreement about a particular policy and the length of time such a policy will take to be passed. Golub (2007: 157) links this to the notion of policy stability, noting that ‘as it becomes numerically easier to form a winning coalition, policy becomes less stable and easier to change’. König (2008) elaborates on this connection, noting that the time-consuming side-payments and compromises required to achieve policy change in a heterogeneous preference environment are unnecessary if there is broad agreement, and that ‘this should greatly speed up decision-making’ (see also König, 2007: 420–1). In both cases, the authors relate policy agreement directly to the level of the hazard of passage, with greater heterogeneity and a larger ‘core’ resulting in greater policy stability and correspondingly lower hazards of passage.

An alternative interpretation, however, might consider how heterogeneity and policy stability might influence the process of passage, and thereby consider their effects on the shape as well as the overall level of the hazard. Golub (2007: 157), for example, argues that, ‘[i]n the extreme case, stability precludes the adoption of new legislation regardless of how long the actors negotiate’; such a position suggests that, for such an attempted policy change, the likelihood that a compromise will be reached (and therefore that passage will occur) declines over the course of the decision-making process. Conversely, a lack of such stability might suggest that – perhaps in addition to hazards of passage being higher – the likelihood that policy change would occur is increasing (or, at least, decreasing at a slower rate) in the length of the decision-making process. In other words, one could hypothesize that the slope of the hazard of adoption is decreasing in the degree of policy stability, as measured by the indicators common to the literature.⁷

The expectation offered here is but one of a number of possible hypotheses one might derive about how exogenous influences shape the hazard of EU policy adoption. Moreover, the example is a relatively simple one; among other shortcomings, it fails to account for the mediating role of institutions and institutional change (including such significant variables as qualified majority voting) in the relationship between member preferences, interest heterogeneity and EU decision-making. At the same time, the example illustrates how one might begin to think theoretically about the dynamics of

the decision-making process itself, and in particular about whether and how the institutional and strategic conditions in the European Union alter the prospects for policy change over time.

Conclusion

In his introduction to a recent study, König (2007: 417) correctly notes that ‘rich and systematic data analyses on the process of EU legislative decision making can help answer the most important questions of EU scholars’. By drawing attention to and integrating a range of conceptual, empirical and statistical issues, the studies in this and the forthcoming issue advance that project in numerous ways, and their conclusions highlight the substantial work remaining on this important topic. Golub and Steunenberg’s paper, in particular, underscores the potential for extracting greater substantive understanding out of quantitative analyses, through increased sensitivity to how the causal dynamics of the EU decision-making process change over time.⁸

At the same time, however, it is also important to understand the limits of our models. Empirically speaking, the occurrence of covariate effects that change over time is consistent with a range of possible data-generating processes. Here, I have focused on one: the case when the ‘shape’ of the baseline hazard – that is, the conditional probability of the event of interest – itself varies systematically with observable characteristics of the data. Such a situation deserves attention, for at least two reasons. First, as illustrated here, it can mimic the more classic cause of non-proportionality, that of changing covariate influences over (process) time. Second, the awareness that the shape of a hazard can differ across different institutional conditions or constellations of member preferences should provoke additional theorizing about why such differences might exist.

Such potential innovations are not without their challenges. Foremost among these is the need for careful theorizing; parametric approaches that attempt to model the shape of the baseline hazard require a compelling theoretical basis, and a key comparative strength of the Cox approach is its relative flexibility and robustness in the absence of such theory. Moreover, in thinking about how and why hazards change over time, a key factor for consideration is *heterogeneity*: the degree to which, conditional on a model’s covariates, observations in the data are not strictly exchangeable. It is well understood that, in the context of survival data, heterogeneity generally leads to declining hazards; moreover, disentangling ‘substantively interesting’ influences on the shape of the hazard from those due to heterogeneity is not always easy (see e.g. Zorn, 2000). At the same time, the potential for gaining real

substantive advances in our understanding of the process of EU decision-making from that approach is great, making such a project worthy of investigation.

Notes

- 1 See Seal (1977) for an overview of the early development of this literature.
- 2 See Denson (1998) and Oakes (2001) for reviews of these two disparate histories.
- 3 The two periods are often separated by a 'washout' period, in which the effects of the treatment/placebo from period 1 are allowed to wash out and the subject's condition to return to its original state.
- 4 See Feingold and Gillespie (1996) for a discussion of additional issues relating to the use of crossover designs in time-to-event data.
- 5 Importantly, however, converging or diverging hazards, by themselves, are not necessarily indicative of non-proportionality; in fact, if the hazards for the control group are declining (increasing) over time, then the difference between the treatment and control groups will necessarily have to converge (diverge) for proportionality to be maintained. As equation (4) shows clearly, what is important for the assumption of proportionality is the ratio of the two hazards, not simply whether they are changing over time.
- 6 For simplicity, I use Weibull hazards, of the form $h_i(t) = \lambda_i p(\lambda_i t)^{p-1}$ and a single binary covariate X . In panel (a), the hazards have $\lambda_i = \exp(-1X_i)$ and $p = .75$; in panel (b), the hazard has $\lambda_i = \exp\{-1[X_i \times \ln(t)]\}$ and $p = .75$. In panel (c), X is unrelated to λ (that is, its coefficient is 0) but p is defined as $p_i = \exp[-0.28768 + (-0.5X_i)]$. In this way, the shape of the 'baseline' (i.e. $X = 0$) hazard is identical in all three panels, and all of the relevant variation in the dashed ($X = 1$) hazard is owing to X 's influence on λ and/or p .
- 7 Note too that such an approach might provide a means of testing the constructivist expectation of convergence. For example, a finding that – even among actors with otherwise divergent preferences – the hazard of policy adoption is nonetheless increasing in the length of deliberations would be consistent with the idea that actors' (constructed) preferences converge over time. Moreover, if the extent of that convergence (measured as the degree to which hazards were rising over time) was itself increasing in conventional measures of policy heterogeneity, the support for the convergence proposition would be even stronger.
- 8 I should note as well that their recommendations regarding time-varying covariates and non-proportionality remain important, even if one is not interested in directly modeling the shape of the underlying hazard.

References

- Armitage, Peter and Michael Hills (1982) 'The Two-Period Crossover Trial', *The Statistician* 31(2): 119–31.

- Cox, David R. (1972) 'Regression Models and Life Tables (with Discussion)', *Journal of the Royal Statistical Society, Series B* 34: 187–220.
- Denson, William (1998) 'The History of Reliability Prediction', *IEEE Transactions on Reliability* 47(3SP): 321–8.
- Feingold, Marcia and Brenda Wilson Gillespie (1996) 'Cross-Over Trials with Censored Data', *Statistics in Medicine* 15: 953–7.
- Golub, Jonathan (2007) 'Survival Analysis and European Union Decision-Making', *European Union Politics* 8(2): 155–79.
- Jones, Byron and Michael G. Kenward (2003) *Design and Analysis of Cross-Over Trials*, 2nd edn. London: Chapman & Hall.
- König, Thomas (2007) 'Divergence or Convergence? From Ever-Growing to Ever-Slowing European Legislative Decision Making', *European Journal of Political Research* 46 (May): 417–44.
- König, Thomas (2008) 'Analysing the Process of EU Legislative Decision-Making: To Make a Long Story Short . . .', *European Union Politics* 9(1), forthcoming.
- Oakes, David (2001) 'Biometrika Centenary: Survival Analysis', *Biometrika* 88(1): 99–142.
- Seal, Hilary (1977) 'Studies in the History of Probability and Statistics. XXXV: Multiple Decrements or Competing Risks', *Biometrika* 64(3): 429–39.
- Tsebelis, George (2002) *Veto Players: How Political Institutions Work*. Princeton, NJ: Princeton University Press.
- Zorn, Christopher (2000) 'Modeling Duration Dependence', *Political Analysis* 8 (Autumn): 367–80.

About the author

Christopher Zorn is Professor, Department of Political Science,
Pennsylvania State University, University Park, PA 16802, USA.

Fax: +1 814 863 8979

E-mail: zorn@psu.edu
