This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Age-Period-Cohort Analysis: a Summary of Analytical Approaches and Results

Enrico Debiasi – enrico.debiasi@ekh.lu.se

Centre for Economic Demography & Department of Economic History

Lund University – January 2018

Introduction

The importance of distinguishing the effects of age, period, and cohort is given by the different nature of the relationship that these time components have with the outcome of interest. Age effects are internal to the individual and they reflect the biological and social processes of aging (Yang & Land, 2013). Period effects arise from events and changes happening as time passes by that affect individuals of all ages, for example: wars, famine, policy changes. Finally, cohort effects derive from differences between groups of people who go through a common initial event (e.g. birth) in the same time unit (e.g. year). Cohort effects arise from a variety of time related changes: firstly, the similar experience that birth cohorts have in going through historical and social event at the same age, thus indicating the intersection of individual level characteristics and macrosocial influences; secondly, birth cohorts continuously change the composition of the population thus reflecting social change (Ryder, 1965). As a whole, APC analysis allows to describe the complex social, historical and environmental factors that simultaneously impact individuals and populations (Yang & Land, 2013).

This paper aims at summarizing methods for age-period-cohort analysis reporting their main characteristics as well as advantages, disadvantages and related relevant results. Eventually it provides an organized classification of methodologies and important literature which is useful for researchers interested in approaching this type of analysis.

The LONGPOP project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Theoretical concepts and the identification problem

The three time components at the basis of demographic (but also epidemiologic and other social sciences) analysis are age (A), period (P), and cohort (C). They highlight different perspectives when interpreting time related changes in the outcome of interest. Briefly, *age effects* are related to the aging process of individuals and therefore are "the effects of differences in the ages of the individuals at the time of observation on an outcome of interest" (Yang & Land, 2013). *Period effects* are those related to events that occur at a specific point in time and affect all people of all ages " the effects of differences in the time periods of observation or measurement of the outcome" (Yang & Land, 2013). *Cohort effects*, instead, are those arising from characteristics that are shared from a group of individuals going through the same event in the same span of time. The most common groups of this kind found in scientific studies are birth year cohorts, meaning, for example, all the people born in a certain year (but it could also be, for instance, all people that married in a certain month). In other words, "the effects of differences in the year of birth or some other shared life events for a set of individuals" (Yang & Land, 2013). It may, sometimes, be confusing to distinguish these three effects on a conceptual level and only providing formal definitions might not be the best way to clarify the nuances. Suzuki (2012) gives an informal way to understand and to place them in the everyday life; the author reports a fictional conversation between two workers:

Senior worker: "I can't seem to shake off this tired feeling. Guess I'm just getting old." [Age effect]

Junior worker: "Do you think it's stress? Business is down this year, and you've let your fatigue build up." [Period effect]

Senior worker: "Maybe. What about you?"

Junior worker: "Actually, I'm exhausted too! My body feels really heavy."

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Senior worker: "You're kidding. You're still young. I could work all day long when I was your age."

Junior worker: "Oh, really?"

Senior worker: "Yeah, young people these days are quick to whine. We were not like that." [Cohort effect]

Arguably, the main point of discussion about age-period-cohort analysis in the literature is the *identification problem*. The identification problem arises from the impossibility of disentangling age, period, and cohort effects in a unique way given their perfect multicollinearity: by knowing two components it is possible to find the third by linearly combining them. For example:

$$age = period - cohort$$

It follows that ordinary least square or maximum likelihood estimators of generalized linear models (such as linear regression, log-linear regression, logistic models) do not exist. The coefficients cannot be uniquely defined as there is an infinite number of possible solutions. Hence, it is not possible to estimate age, period, and cohort effects unless one or more constraints on the coefficients are imposed (a mathematical outline of the identification problem can be found in Yang & Land (2013)).

Analytical approaches to APC analysis and to the identification problem

Since the first considerations on the concept of cohort effects, followed by the analysis of time effects in more and more detail, arriving to today's sophisticated age-period-cohort model, there has been a lively debate on the pro and cons of each method and on assessing which way leads to the most reliable results. Many propositions have been put forward creating a vast literature on the subject; it is therefore relevant to

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

organize such contributions in terms of timing and features in order to make an informed choice when approaching suitable data.

One of the first contribution is the seminal work by Frost (1939) who pointed out the interesting perspectives that can be studied when taking into account all three components. His approach was descriptive and he mainly used graphs to display his theory; he studied data about tuberculosis in Massachusetts at the turn of the twentieth century and showed how differences in mortality for certain age groups changed through time and by cohorts and how age-specific death rate showed a more regular pattern when looked from cohorts than from periods. His study highlighted the importance of early life conditions compared to current conditions on the development of a disease that has long latency.

There has been an increasing interest in writing and publishing scientific works related to age-period-cohort analysis and findings (Figure 1).

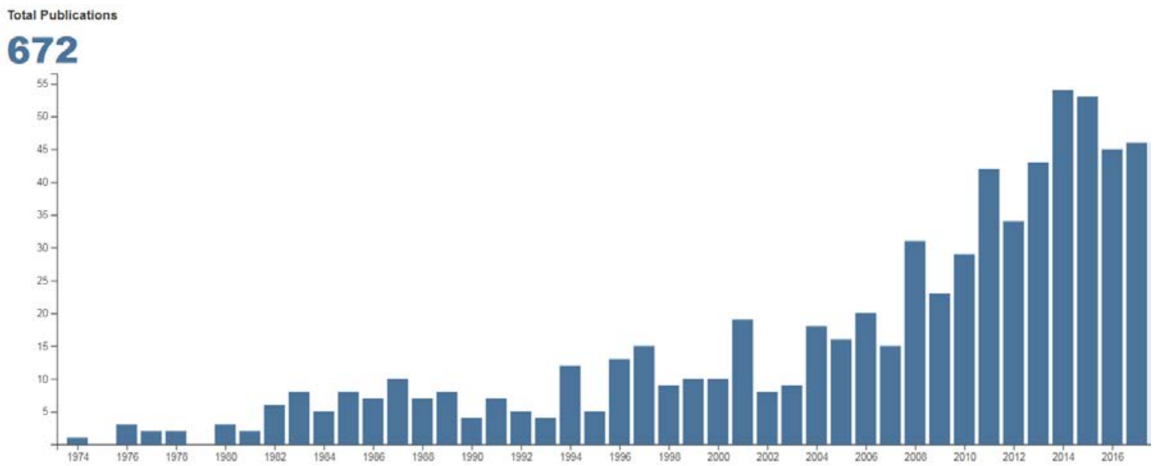

Figure 1: Number of publications within the Web of Science database that have been written in English and that contain the words "age period cohort" in the title. (https://apps.webofknowledge.com)

In time, there has been a shift from descriptive studies to more and more empirical methods. And with this shift the attention has moved to the identification problem. Several authors have presented different categorization of age-period-cohort models (Holford, 1991; Robertson, Gandini, & Boyle, 1999; Yang & Land, 2013). This review uses this references as a starting point to provide a complete outline and classification of the models used in the last 60/70 years.

Descriptive approaches

Chronologically, the first analyses discussing age, period, and cohort in relation to each other have been of a descriptive nature and used indicators such as *age standardized rates* or *age specific rates*. Such measures are still in use today's studies as they are often useful to present the data and provide insightful graphical description of time trends.

Age standardized rates are the result of the weighted average of age specific rates. They are usually reported per 100 000 persons and the weight typically refers to the proportion of the population in the age group. Figure 2 presents an example of such indicator and displays the age standardized death rates by cause of death in the Unites States between 1970 and 2002 where the population in 2000 was used a standard population.

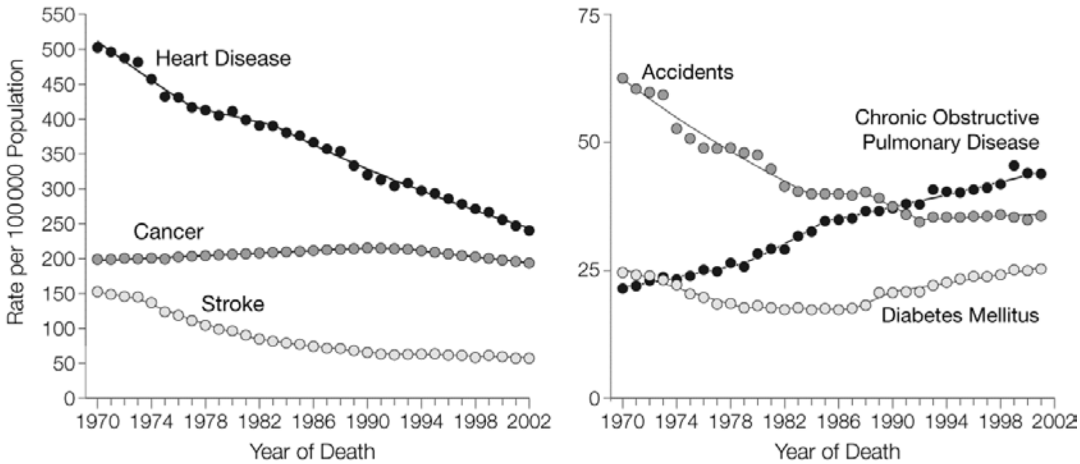

Figure 2: Age standardized death rates by cause of death in the Unites States between 1970 and 2002. 2000 US population as standard. (Jemal, Ward, Hao, & Thun, 2005)

The drawback of such indicator is that the rate depends on the period chosen as standard as the results will be based on the age composition of it. This means that trends over a medium/long time span might not be reliable because the age composition has varied sensibly through the years. The population pyramid in Figure 3 clearly show how the population has aged in the last 60+ years from 1950 to 2017.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Figure 3: On the left, population pyramid for the United States in the 1950; on the right for the year 2017. (www.populationpyramid.net)

Alternatively, studies may report age specific rates plotted for different periods or cohorts on the same graph displaying the evolution of the outcome through the time component of interest. Figure 4 displays, on the left, the age specific death rate from bladder cancer for different cohorts of men aged 40 to 69 in England and Wales between 1951 and 1980. The right hand side of Figure 3 shows the age specific mortality rate from lung cancer for women aged 25 to 69 in the same area in the same period.

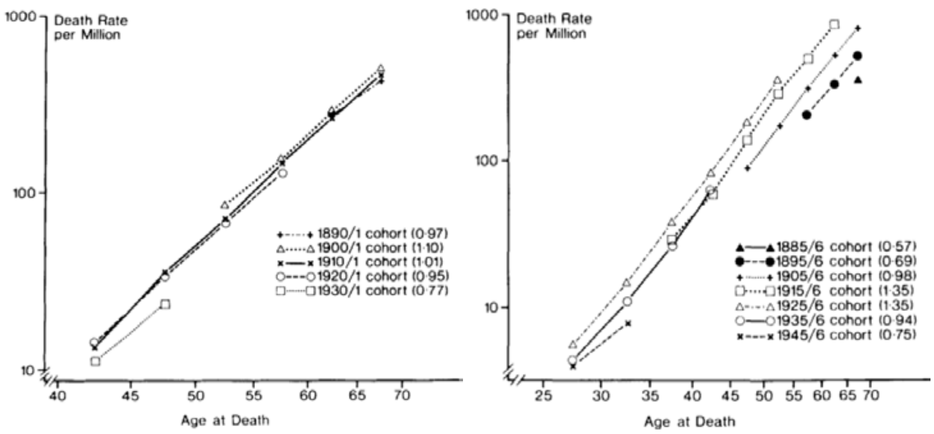

Figure 4: the left panel displays the age specific death rate from bladder cancer for different cohorts of men aged 40 to 69 in England and Wales between 1951 and 1980. On the right the age specific mortality rate from lung cancer for women aged 25 to 69 in England and Wales between 1951 and 1980. (C. Osmond & Gardner, 1982)

The drawbacks of these type of graph is that, as shown in Figure 5, there might be inconsistencies between age-period and age-cohort plots: the graph reports incidence

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

rates for lung cancer in women aged 20 to 84 who lived in Connecticut between 1940 and 1984. The solid line reports the age specific rate by period whereas the dotted line indicates the age specific incidence rate by cohort. While the trend showed by the cohort lines is what could be epidemiologically expected (a rate increasing with age throughout all cohorts), the trend reported by periods contradicts this expectation showing a flattening after the age of 50 (Yang & Land, 2013).

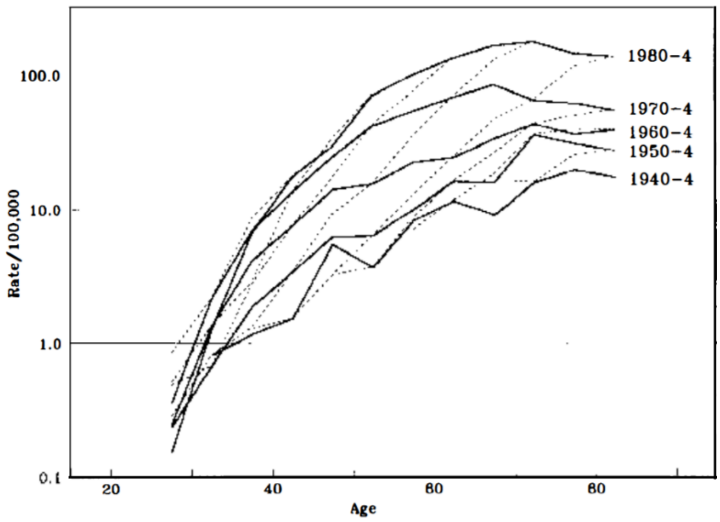

Figure 5: lung cancer incidence rates for women in Connecticut by age. The solid lines are for constant periods, while dotted lines are for constant cohorts. (Holford, 1991)

It is therefore possible to understand why these two-way age by period or age by cohort graphs are useful only for descriptive purposes and not to quantitatively assess the source of change and how the three effect operates (Kupper, Janis, Karmous, & Greenberg, 1985). In other words, even though these indicators are a good way to introduce time trends, neither one of them represent a full APC model and consequently it is not possible to obtain age, period, and cohort effects. In addition, as summarized by Osmond (1985), these two measures present an important disadvantage that make them unsuited for detailed investigations: it is possible to depict only two of the three time component simultaneously, so there will be either rates by period or by cohort.

Given the limitations of descriptive analysis in understanding the different mechanisms for the three time component, it is clear that there is the need for more sophisticated statistical modelling. Historically the evolution of age-period-cohort analysis from visual approaches using graphs to more qualitative estimation of the effects in play, have also been favored by the development of computing power and by the refinement of statistical software which nowadays include default packages dedicated to APC studies.

Reduced two factors models

One way of dealing with the identification issue is to use a reduced two-factor model. The idea behind this method is to include only one or two time components in the model specification (usually age and period):

$$\log(E_{ap}) = \log\left(\frac{I_{ap}}{P_{ap}}\right) + \mu + \alpha_a + \beta_p$$

where E_{ap} is the expected incidence rate, I_{ap} is the number of cases and P_{ap} is the total population at risk in the age-period cell. Additionally, μ is the intercepts α_a is the age effect and β_p is the period effect. When considering only age and period effects, it is assumed that changes in rates over time are equal across cohorts and that period effects, such as the economic situation or technological changes independent from the cohort, play the major role in explaining these changes. However, this assumption disagrees with an important part of the literature arguing for the importance of cohort effects (Barker, 1998; Finch & Crimmins, 2004; Fogel & Costa, 1997; Ryder, 1965 and see Fogel, 2003)

Kupper et al. (1985) explain further limitations of the above approach; they note that besides not including cohort effects, it does not contain any interaction term between age and period which could be interpreted as cohort effects. In other words, the expected value for the age-period (a, p) cell of a table reporting rates by age and period groups is only determined by the marginal effects of the ath row (age) and pth column (period), and it does not take into consideration a possible joint effect of the two (cell specific). While the exclusion of an interaction term introduces the need for an additional assumption, the authors graphically show that it could be difficult to find support for it.

LUND UNIVERSITY

LONGPOP

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

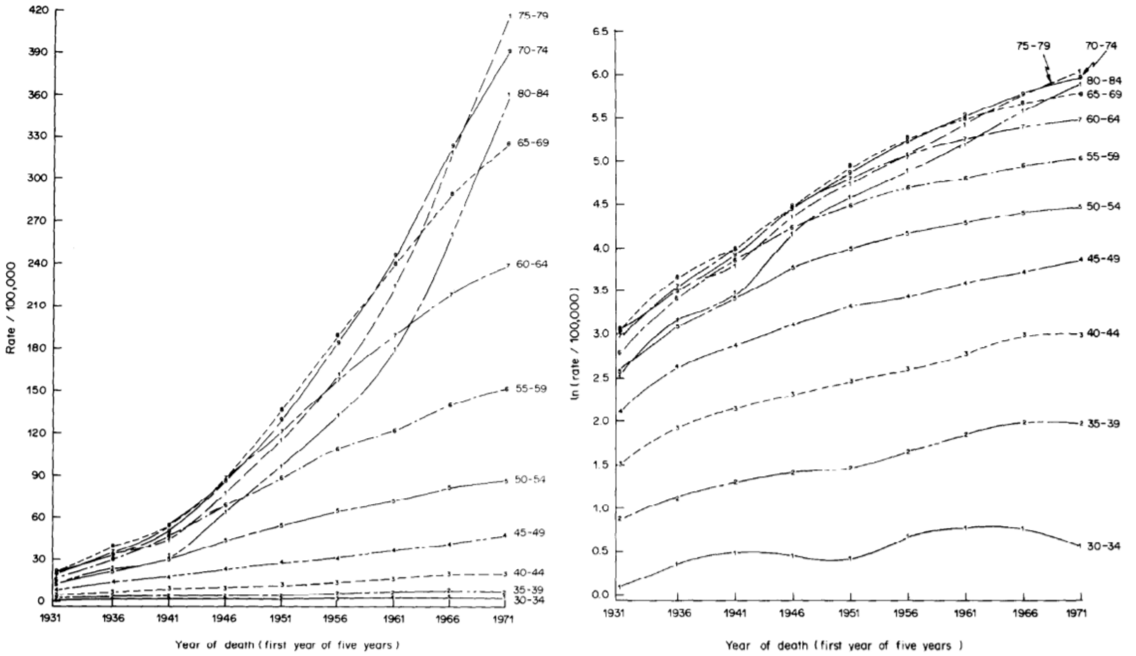

Figure 6: On the left, age specific rate per 100,000 by year of death for different age groups. On the right, rates are reported on a logarithmic scale. Adapted from (Kupper et al., 1985)

Kupper and colleagues (1985) argue that a lack of parallelism in the graph reporting age-specific rates by period (both on a linear and log scale) is evidence for the presence of a cohort effect. The theoretical and graphical lack of support for the underlying assumption of such method opens the discussion to other approaches.

In summing up the two factors approaches, the tactic has been to assess the effect of either two of the three components while neglecting the effect of the remaining one (or arguing for a negligible importance of the excluded factor). However, age, period, and cohort effects are often linked to different causal interpretations and therefore not controlling for one of these three components may lead to spurious effects (K. O. Mason, Mason, Winsborough, & Poole, 1973).

Constrained Generalized Linear Models

There are essentially two ways to constrain the coefficients in order to avoid the identification problem: just-identified constraints and over identified constraints. Just identified models are built by adding to the three time effect a restriction such as forcing

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

two age groups to have the same effect (W. M. Mason & Smith, 1985). The second type – over identified constraints – force an additional restriction on the model by changing the size of age, period, and cohort groups. Usually this is achieved by having larger, less refined, groups for periods or cohorts and smaller age categories (Yang & Land, 2013).

The main disadvantage common to Constrained Generalized Linear Models is that in applying them, there is the need to rely on additional information from external sources in order to find reasonable constraints, however it is often the case that such information is not available. On a similar line of argument, the findings of studies adopting this approach become completely dependent on which constraint is chosen and consequently different results are found when changing the identification restriction and, in case of over identification, on the chosen size for age, period, and cohort groups (W. M. Mason & Smith, 1985; Yang, Fu, & Land, 2004).

Proxy variables

The idea behind the proxy variable approach is straightforward: as the name suggests, one or more proxy variable is used to replace the age, period, or cohort variable in the model. While relative cohort size can be used as a proxy for cohort effect (O'Brien, Stockard, & Isaacson, 1999), unemployment rate or labour force size might be used to substitute period effects (Pavalko, Gong, & Long, 2007). Usually either period or cohort proxies are used. When the cohort variable is replaced the model is called age-period-cohort characteristics model (APCC) (O'Brien et al., 1999). This method is particularly interesting if we consider the measure of age, period, and cohort effects as an unsatisfactory measure of demographic changes:

> *"The phrase "age, period, and cohort effects" is probably an unfortunate one. Ages, periods, and cohorts do not have either direct or indirect effects on demographic or social phenomena. Age is a surrogate--probably a very good one in most applications--for aging or more generally for physiological states, amount of exposure to certain social influences, or exposure to social norms. […] However, it is clear that individuals' age physiologically and socially at different rates. "Period" is a poor proxy for some set of contemporaneous influences, and*

"cohort" is an equally poor proxy for influences in the past. (Hobcraft, Menken, & Preston, 1982)."

From this is follows that whenever possible, to understand the sources of variation in vital rates, it is better to use the underlying variables for which age, period, and cohort factors are proxies (Hobcraft et al., 1982).

Nevertheless, there are some drawbacks in using proxy variables: in first instance, the variables chosen as proxies should not be linearly related to the factor they substitute. If that is the case the multicollinearity would still be in place. Secondly, somewhat in relation to the above Hobcraft and colleagues quote, the substitution of age, period, and cohort with measured variables does not necessarily lead to a better model. Whilst it might solve the identification problem, it opens up to the possibility of having an incorrectly specified model. In other words, if the chosen proxy variables do not account for the full variation of the factors they aim at substituting, such approach will be insufficient to obtain age, period, and cohort effects (Smith, Mason, & Fienberg, 1982).

Penalty function approach

This approach can be seen as an alternative way to define model constraints. The penalty function is a measure of the distance between the two factors models (age-period, age-cohort, and period-cohort) and the three factors model (age-period-cohort). Robertson et al. (1999) report a summary of the steps involved in this type of estimation:

- firstly, the age-period, age-cohort, and period-cohort models must be estimated and the parameters obtained saved;
- then, the full age-period-cohort model must also be estimated by imposing a constrain on the parameters that could be arbitrary (e.g. make the effects of the first and last period equal) or based on a priori knowledge (biological hypothesis based on previous knowledge of the event under study) (see Decarli & La Vecchia, 1987);
- thirdly, the penalty function is derived as the sum of the squares of the differences between the parameters of each of the three two-factor models and

the full three factor model weighted by a measure of goodness of fit of each of the three two-factor models, such as the deviance.

- Finally, the value obtained from the minimization of the penalty function is used together with the parameters of the full model to get identifiable estimates of age, period, and cohort.

Other works based on this approach are Barrett (1973) and Osmond & Gardner (1982) who use arbitrary constraints and from(Fienberg & Mason, 1979; K. O. Mason et al., 1973) who constrain the parameters based on prior knowledge (Decarli & La Vecchia, 1987). More recent contributions are Fu (1998) and Fu (2000).

Nonlinear parametric transformation

The identification problem can be overcome by including a nonlinear function of at least one of the time factors in the model. A classic example for this type of models is the inclusion of age squared to explain the nonlinear relation with the outcome of interest.

This approach has however two main disadvantages: firstly, it might be difficult to decide which nonlinear function should be included in the specification; secondly, and perhaps even more important, is the fact that the use of this method does not lead to a complete solution of the identification problem because the linear effects remain undefined.

Individual record approach

Another way to overcome the indefinability problem was proposed by Robertson & Boyle (1986) who suggested that, with the use of individual records, the linear relationship between age, period, and cohort could be broken. Their idea was that with more detailed data it is possible to assign to each age-period group two distinct cohorts.
If in a usual two way age (rows)-period(columns) table each cell (cohort) can be identified by $k=j-i+I$ where $i = (1,\dots,I)$ indicates the age groups, $j = (1,\dots,J)$ indicates the period groups and $k = (1,\dots,K)$ indicates cohort groups, by dividing each cell diagonally the lower triangle is now identified with $k=j-i+I$ and the upper triangle is now identified with $k=j-i+I+1$. In this way, the age and the period groups do not point to a unique cohort solving the identifiability issue. Graphically:

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

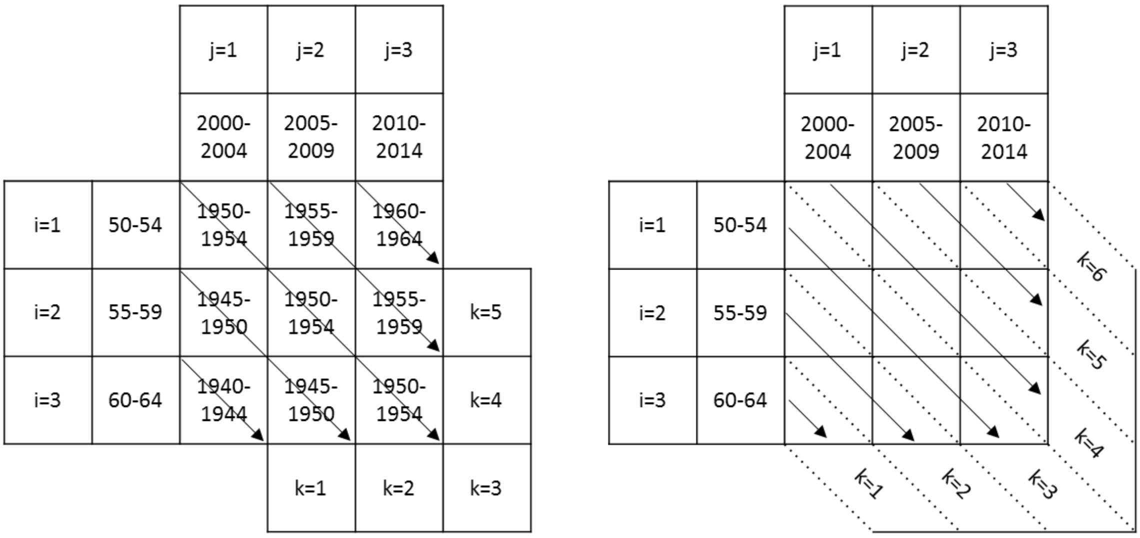

Figure 7: on the left, the table shows the situation when using aggregate data; for each age-period cell there is only one cohort related to it. When using individual level data, it is possible to divide people in two cohort groups for each age-period cell.

$$k = j - i + I$$

$k = j - i + I$ for the lower triangle

$k = j - i + I + 1$ for the upper triangle

As Figure 7 shows, in the right column, the square is divided into two triangles and this allows for the identifiability problem to be solved (Robertson & Boyle, 1986).

However, as Osmond & Gardner (1989) noted, the use of individual records does not come without problems. They observe that there is still the need for underlying assumptions that sometimes might not be justifiable. The fact that the use of individual record is not enough to come to a strong solution suggest that the APC is not a problem of data but a problem of method (Fu, 2008).

<u>Splines and stastical packages</u>

The use of splines to deal with the identification issue has gained attention following the work of Bendix Carstensen (2007; 2006) who developed functions for the statistical software R that fit age-period-cohort models and to graphically display results.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

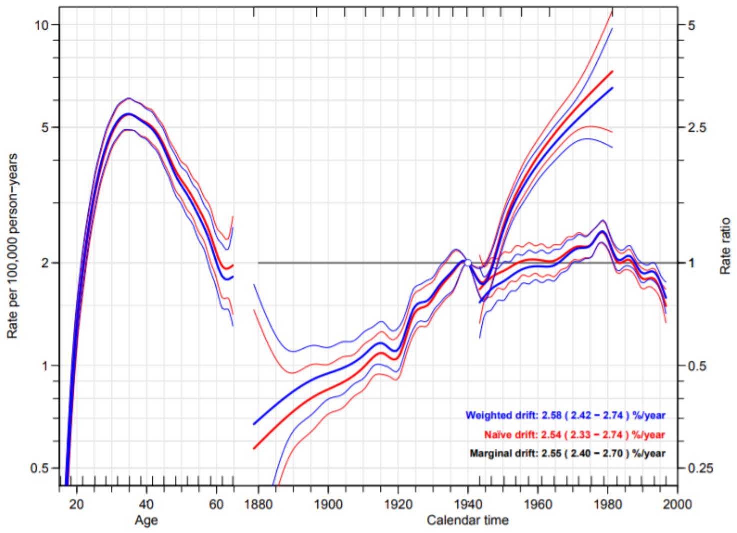

Figure 8: Graphical display of the results obtained with the R package designed by Carstensen (2006, 2007). Age effects are showed on the left and are reported as incidence rate. Cohort and period effects are the curves in the center and on the right of the graph respectively; they are reported as rate ratios with respect to the reference cohort.

Additional packages have been written for the statistical software STATA first by Rutherford, Lambert, & Thompson (2010) and then Sasieni, (2012) proposed a complementary approach. The first authors wrote the **apcfit** command that basically translates Carstensen R package in STATA and they additionally provide the **poprisktime** command that helps in transforming the dataset in the proper form to run the age period cohort analysis using **apcfit**.

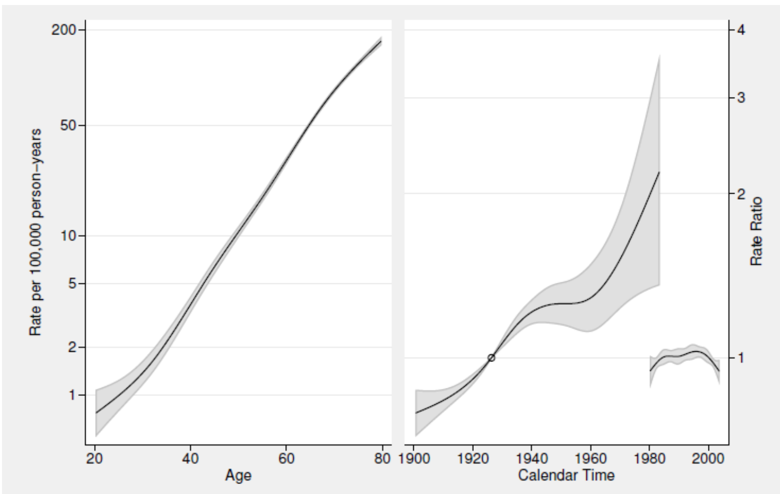

*Figure 9: graph resulted from the application of the **apcfit** STATA package for age-period-cohort analysis. From left to right, the chart reports age, cohort, and period effects. Age effects are displayed as incidence rate whereas cohort and period effect are showed as rate ratio with respect to the reference cohort (highlighted with the circle). The shaded area represents the 95% confidence interval.*

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Sasieni further developed the potentiality of age period cohort analysis in STATA with the command **apcspline** that facilitate the extrapolation of the model fit for making future projections.

Generalized Linear Mixed Models

One of the most recent contributions in terms of methodologies have been developed in a Generalized Linear Mixed Models (GLMM) context. Such framework extends the analytical potential of the Generalize Linear Models by allowing to specify both fixed and random parameters at the same time. Within this setting, Yang and Land developed the class of Hierarchical Age Period Cohort (HAPC) models (Yang, 2006; Yang & Land, 2006, 2008). The underlying idea of these models is that time periods and cohort membership represent the social historical context and individuals are embedded in this context. This conceptualization is then translated in the model by specifying age as a fixed effect, and period and cohort as random effects (Yang & Land, 2013). A clear example regarding the application of such methodology is given by Master and colleagues (2012). One of the advantage, evident from their work is the easy interpretation of the three time effects, that can be reported as death rates, through a set of separate graph for the age, period, and cohort impact on the outcome as shown in Figure 10.

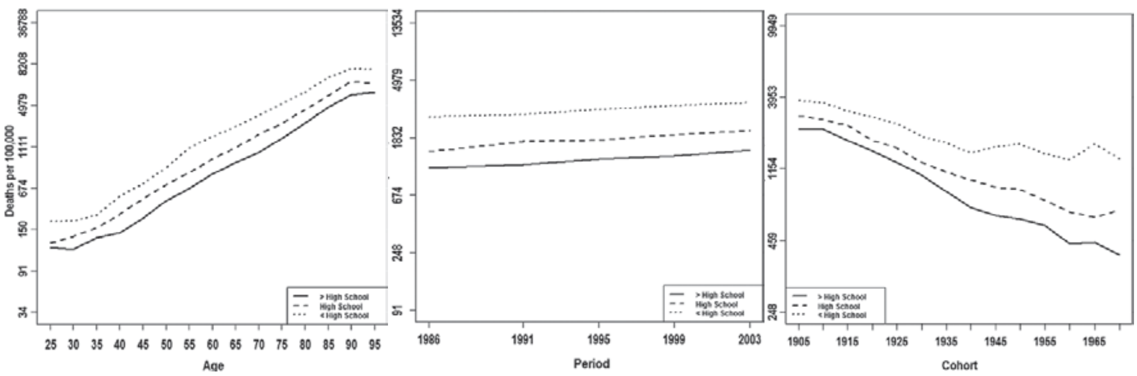

Figure 10: Age-period-cohort effects from a HAPC model (adapted from Masters et al., 2012)

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Concluding remarks

In the last decades age-period-cohort studies have been at the center of a constructive debate among authors interested in investigating and explaining time effects more and more in detail. This article outlined in an organized manner the main methods produced from such branch of the literature, highlighting pro and cons of each methodology.

From the first studies to nowadays investigations there has been a shift from descriptive approaches to methods able to disentangle the three different effects of age, period, and cohort more clearly. Nevertheless, these methodologies rely on sometimes strong assumptions and therefore they must be applied carefully. Finally, while there is a good selection of methods for aggregated data, there is a need for ways to better and fully exploit the increasing availability of large, individual level datasets.

References

Barker, D. J. (1998). In utero programming of chronic disease. *Clinical Science (London, England: 1979)*, *95*(2), 115–128.

Barrett, J. C. (1973). Age, time and cohort factors in mortality from cancer of the cervix. *The Journal of Hygiene*, *71*(2), 253–259.

Carstensen, B. (2007). Age-period-cohort models for the Lexis diagram. *Statistics in Medicine*, *26*(15), 3018–3045. https://doi.org/10.1002/sim.2764

Carstensen, Bendix. (2006). *Demography and epidemiology: Age-Period-Cohort models in the computer age*. Department of Biostatistics, University of Copenhagen.

Decarli, A., & La Vecchia, C. (1987). Age, period and cohort models: review of knowledge and implementation in GLIM. *RIVISTA DI STATISTICA APPLICATA*, *20*(4), 397–410.

Fienberg, S. E., & Mason, W. M. (1979). Identification and Estimation of Age-Period-Cohort Models in the Analysis of Discrete Archival Data. *Sociological Methodology*, *10*, 1–67. https://doi.org/10.2307/270764

Finch, C. E., & Crimmins, E. M. (2004). Inflammatory exposure and historical changes in human life-spans. *Science (New York, N.Y.)*, *305*(5691), 1736–1739. https://doi.org/10.1126/science.1092556

Fogel, R. W. (2003). *Changes in the Process of Aging During the Twentieth Century: Findings and Procedures of the Early Indicators Project* (Working Paper No. 9941). National Bureau of Economic Research. https://doi.org/10.3386/w9941

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Fogel, R. W., & Costa, D. L. (1997). A Theory of Technophysio Evolution, With Some Implications for Forecasting Population, Health Care Costs, and Pension Costs. *Demography*, *34*(1), 49–66. https://doi.org/10.2307/2061659

Frost, W. H. (1939). The Age Selection of Mortality from Tuberculosis in Successive Decades. *American Journal of Hygiene*, *30*, 91–96.

Fu, W. J. (1998). Penalized Regressions: The Bridge versus the Lasso. *Journal of Computational and Graphical Statistics*, *7*(3), 397–416. https://doi.org/10.2307/1390712

Fu, W. J. (2000). Ridge estimator in singulah oesiun with application to age-period-cohort analysis of disease rates. *Communications in Statistics - Theory and Methods*, *29*(2), 263–278. https://doi.org/10.1080/03610920008832483

Fu, W. J. (2008). A Smoothing Cohort Model in Age–Period–Cohort Analysis With Applications to Homicide Arrest Rates and Lung Cancer Mortality Rates. *Sociological Methods & Research*, *36*(3), 327–361. https://doi.org/10.1177/0049124107310637

Hobcraft, J., Menken, J., & Preston, S. (1982). Age, Period, and Cohort Effects in Demography: A Review. *Population Index*, *48*(1), 4–43. https://doi.org/10.2307/2736356

Holford, T. R. (1991). Understanding the Effects of Age, Period, and Cohort on Incidence and Mortality Rates. *Annual Review of Public Health*, *12*(1), 425–457. https://doi.org/10.1146/annurev.pu.12.050191.002233

Jemal, A., Ward, E., Hao, Y., & Thun, M. (2005). Trends in the Leading Causes of Death in the United States, 1970-2002. *JAMA*, *294*(10), 1255–1259. https://doi.org/10.1001/jama.294.10.1255

Kupper, L. L., Janis, J. M., Karmous, A., & Greenberg, B. G. (1985). Statistical age-period-cohort analysis: A review and critique. *Journal of Chronic Diseases*, *38*(10), 811–830. https://doi.org/10.1016/0021-9681(85)90105-5

Mason, K. O., Mason, W. M., Winsborough, H. H., & Poole, W. K. (1973). Some Methodological Issues in Cohort Analysis of Archival Data. *American Sociological Review*, *38*(2), 242–258. https://doi.org/10.2307/2094398

Mason, W. M., & Smith, H. L. (1985). Age-Period-Cohort Analysis and the Study of Deaths from Pulmonary Tuberculosis. In *Cohort Analysis in Social Research* (pp. 151–227). Springer, New York, NY. https://doi.org/10.1007/978-1-4613-8536-3_6

Masters, R. K., Hummer, R. A., & Powers, D. A. (2012). Educational Differences in U.S. Adult Mortality: A Cohort Perspective. *American Sociological Review*, *77*(4), 548–572. https://doi.org/10.1177/0003122412451019

O'Brien, R. M., Stockard, J., & Isaacson, L. (1999). The enduring effects of cohort characteristics on age-specific homicide rates, 1960-1995. *AJS; American Journal of Sociology*, *104*(4), 1061–1095.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 676060

Osmond, C., & Gardner, M. J. (1982). Age, period and cohort models applied to cancer mortality rates. *Statistics in Medicine*, *1*(3), 245–259. https://doi.org/10.1002/sim.4780010306

Osmond, C., & Gardner, M. J. (1989). Age, period, and cohort models. Non-overlapping cohorts don't resolve the identification problem. *American Journal of Epidemiology*, *129*(1), 31–35.

Osmond, Clive. (1985). Using Age, Period and Cohort Models to Estimate Future Mortality Rates. *International Journal of Epidemiology*, *14*(1), 124–129. https://doi.org/10.1093/ije/14.1.124

Pavalko, E. K., Gong, F., & Long, J. S. (2007). Women's Work, Cohort Change, and Health. *Journal of Health and Social Behavior*, *48*(4), 352–368. https://doi.org/10.1177/002214650704800402

Robertson, C., & Boyle, P. (1986). Age, period and cohort models: The use of individual records. *Statistics in Medicine*, *5*(5), 527–538. https://doi.org/10.1002/sim.4780050517

Robertson, C., Gandini, S., & Boyle, P. (1999). Age-Period-Cohort Models: A Comparative Study of Available Methodologies. *Journal of Clinical Epidemiology*, *52*(6), 569–583. https://doi.org/10.1016/S0895-4356(99)00033-5

Rutherford, M. J., Lambert, P. C., & Thompson, J. R. (2010). Age-period-cohort modeling. *Stata Journal*, *10*(4), 606–627.

Ryder, N. B. (1965). The Cohort as a Concept in the Study of Social Change. *American Sociological Review*, *30*(6), 843–861. https://doi.org/10.2307/2090964

Sasieni, P. D. (2012). Age-period-cohort models in Stata. *Stata Journal*, *12*(1). Retrieved from http://ageconsearch.umn.edu/record/202123

Smith, H. L., Mason, W. M., & Fienberg, S. E. (1982). Estimable Functions of Age, Period, and Cohort Effects: More Chimeras of the Age-Period-Cohort Accounting Framework: Comment on Rodgers. *American Sociological Review*, *47*(6), 787–793. https://doi.org/10.2307/2095214

Suzuki, E. (2012). Time changes, so do people. *Social Science & Medicine*, *75*(3), 452–456. https://doi.org/10.1016/j.socscimed.2012.03.036

Yang, Y., Fu, W. J., & Land, K. C. (2004). A Methodological Comparison of Age-Period-Cohort Models: The Intrinsic Estimator and Conventional Generalized Linear Models. *Sociological Methodology*, *34*(1), 75–110. https://doi.org/10.1111/j.0081-1750.2004.00148.x

Yang, Y., & Land, K. C. (2013). *Age-Period-Cohort Analysis: New Models, Methods, and Empirical Applications*. CRC Press.