NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

The geometry of mortality change: Convex hulls for demographic analysis

Audrey F. Lai*

Andrew Noymer†

Tsuio Tai‡

23 December 2016

Abstract

We introduce convex hulls as a data visualization and analytic tool for demography. Although heretofore underutilized in population studies, convex hulls are widely used in computer science, and have been applied in academic fields that neighbor demography, such as ecology. We briefly discuss convex hulls, then we show how they may be applied profitably to demography. We do this through three examples, drawn from the relationship between child mortality and adult survivorship (${}_5q_0$ and ${}_{45}p_{15}$ in life table notation). The three examples are: (i) using convex hulls for outlier identification; (ii) for studying sex differences in mortality; and (iii) for studying period and cohort differences. We find, respectively, that convex hulls can be useful in robust compilation of demographic databases, and that the gap/lag framework for sex differences or period/cohort differences is more complex when mortality data are arrayed by two components as opposed to a unidimensional measure such as life expectancy. The potential applicability of these methods goes beyond mortality.

*Department of Computer Science, University of California, Irvine, USA

†Department of Population Health and Disease Prevention, University of California, Irvine, USA. To whom correspondence should be addressed: noymer@uci.edu

‡Department of Sociology, National Taipei University, San Shia, Taiwan

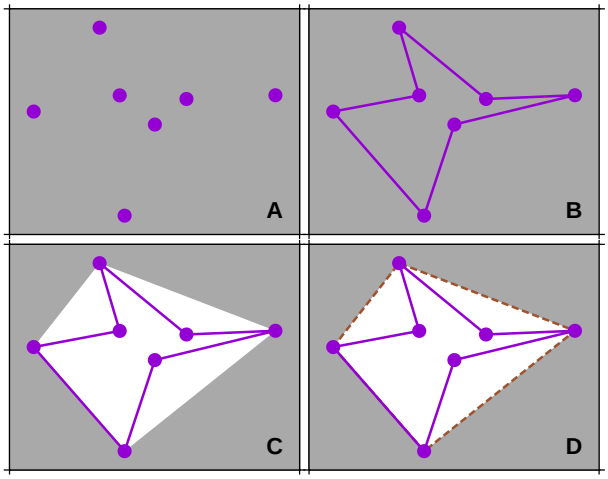


Figure 1: Example. **A**: seven randomly-distributed points in a plane. **B**: Irregular polygon, one possible perimeter that touches all points. **C**: convex hull, shaded in white. **D**: dashed line segments illustrate non-convexity of the original perimeter.

Introduction

We propose convex hulls as a promising technique of demographic analysis, illustrated by three examples. A convex hull of a set of points is a subset of those points, forming a perimeter, in which line segments connecting any two points in the data lie inside or on the hull. An informal heuristic is that if a set of points are pegs in a board, the convex hull is the shape of a rubber band stretched around the outermost pegs, such that all the pegs are enclosed by the band. Convex hulls exist in all dimensions: as a range (line segment) for unidimensional data, as polygons in $\mathbb{R}^2$ ("2D"), as polyhedra in $\mathbb{R}^3$ ("3D"), and as polytopes in higher dimensions. Herein we consider only applications in two dimensions. Figure 1 is an example: the "data" are seven random points in a plane (1A). There are a number of ways to draw a perimeter, of which one is shown (1B). The convex hull, which is unique, is illustrated as a white polygon (1C). The dashed line segments (1D) demonstrate why the perimeter in 1B is not convex. Line segments connecting any two points in the data may be an edge of the convex hull, or interior, but cannot pass outside of it.

Calculating a convex hull of multidimensional data is analogous to sorting unidimensional data, in the sense that it determines the boundaries of the data, which in the univariate case is the minimum and maximum (Barnett 1976). For cross-classified data, the x range is the orthogonal projection of the convex hull onto the x-axis, similarly with the y data, and so on if there are more dimensions. Using convex hulls for data analysis is not a new idea. "Tukey peeling", also called convex peeling (Hodge and Austin, 2004), has some similarities to our first application (below). It entails obtaining robust estimates in multivariate analysis by removing one or more convex hulls from the data, pre-analysis. It dates back to the early 1970s (Huber 1972), and is further elaborated in Tukey (1975) and Bebbington (1978). The properties of convex hulls of data are fairly well understood, provided the data are reasonably well behaved. There is a large literature here; see, f.e., Efron (1965), Eddy (1980), Aldous et al. (1991), Blackwell (1992), Snyder and Steele (1993), Hueter (1994), Massé (2000), Suri et al. (2013). To the best of our knowledge, these techniques have not been applied in depth in demography. Appearances of "convex hull" in the demographic literature are sparse and arise in conjunction with linear programming solutions (f.e., Georgakis and Tziafetas 1982), rather than as a stand-alone analytic tool. Wrigley and Schofield (1981) discuss "demographic terrain" (p. 247), which is similar in spirit to convex hull analysis as we conceptualize it (see also Goldstone 1986, Galloway 1994). Neighboring academic fields such as ecology have used convex hulls more (f.e., Getz and Wilmers 2004, Cornwell et al. 2006). Our three applications illustrate the usefulness of convex hulls to population studies.

Materials and Methods

Using data on all countries in the Human Mortality Database (Barbieri et al. 2015), we analyze the joint distribution of child mortality (${}_5q_0$) and adult survivorship (${}_{45}p_{15}$). The included countries are listed in table 1. Throughout, we refer to cross-classification of child mortality and adult survivorship as the *mortality relationship*, and, as applicable, the *mortality hull*. We perform three analyses, the first of which is to examine outlier countries, in which we systematically delete one country at a time, and quantify how the convex hull changes. The second analysis compares male and female convex hulls, on a per-country basis. The third analysis compares period and cohort

Table 1: **Country list: Start and end years.** Data from Human Mortality Database (2016).

	Period		Cohort	
Country	start	end	start	end
Australia	1921	2011		
Austria	1947	2014		
Belarus	1959	2014		
Belgium†	1841	2012		
Bulgaria	1947	2010		
Canada	1921	2011		
Chile	1992	2005		
Czech Republic	1950	2014		
Denmark	1835	2014	1835	1923
Estonia	1959	2013		
Finland	1878	2012	1878	1921
France*	1816	2014	1816	1923
E Germany	1956	2013		
W Germany	1956	2013		
Greece	1981	2013		
Hungary	1950	2014		
Ireland	1950	2014		
Iceland	1838	2013	1838	1922
Israel	1983	2014		
Italy	1872	2012	1872	1921
Japan	1947	2012		
Latvia	1959	2013		
Lithuania	1959	2013		
Luxemburg	1960	2014		
Netherlands	1850	2012	1850	1921
New Zealand (Maori)	1948	2008		
New Zealand (Non-Maori)	1901	2008		
Norway	1846	2014	1846	1923
Poland	1958	2014		
Portugal	1940	2012		
Russia	1959	2014		
Slovakia	1950	2014		
Slovenia	1983	2014		
Spain	1908	2014		
Sweden	1751	2014	1751	1923
Switzerland	1876	2011	1876	1923
Taiwan	1970	2010		
England and Wales*	1841	2013	1841	1922
Scotland	1855	2013	1855	1922
Northern Ireland	1922	2013		
Ukraine	1959	2013		
United States	1933	2014		

Notes: †: missing 1914–1918.
*: total population (i.e., not only civilian)

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

data, on a per-country and per-sex basis. To calculate the convex hulls, we used IDL ver. 8.6 (Exelis Visual Information Solutions, Inc., Boulder, Colorado, USA). Convex hull calculation is a well-studied problem in computer science (Preparata and Shamos 1985, de Berg et al. 2008), and is available in many software packages.

An example of the convex hull approach is given in figure 2 (p. 6). Here we illustrate the mortality hull, separately by sex, for the entire data set. Individual countries are color-coded, although most of the data are densely clustered and therefore overlapping. Superposed on figure 2 is a more conventional approach, namely a regression fit of ${}_{45}p_{15}$ as a quadratic function of ${}_5q_0$, along with its 95% prediction interval.[1] Near the center of mass, the regression line does a good job of representing the tendency of the mortality relationship. However, we think figure 2 is a good illustration of the shortcomings of the parametric curve fitting approach. Particularly away from the center, even with its prediction interval, the regression line does not represent the variation of the data as well as the convex hull.

The geometric (i.e., convex hull) approach is not meant to be a drop-in replacement for curve fitting, but a complement to it. Nonetheless, in many aspects of population studies, convex hulls better capture the inherent variation of the data, especially in situations where the quantities of interest do not have a homoskedastic relationship.[2] By definition, the shape of the convex hull is determined by extreme values; if this is a shortcoming, it may be avoided by Tukey peeling one or more hulls from the original data prior to the main analysis.

The overarching logic of our approach is that when comparing the mortality relationship (or any other multivariate classification) of two or more populations, convex hulls are a natural way to demonstrate how the data from the countries interleave with each other. This approach is superior to comparing the bivariate ranges, which would replace the hulls with rectangles of potentially much greater area. It is an alternative to fitting curves for each country, and then comparing the two regressions. While convex hulls will not replace the curve fitting approach, it can be a valuable complement, and as shown in figure 2, is quite different.

[1] This is wider than the 95% confidence interval of the regression curve. See, f.e., Snedecor and Cochran (1989), p. 168; DeGroot and Schervish (2002), p. 614; etc.

[2] For quantities analyzed on log scale (f.e., as is often the case with ${}_nM_x$), log the data first, then calculate the hull. This is because, although the logarithmic transformation is monotone, it is not affine, so need it not preserve convexity.

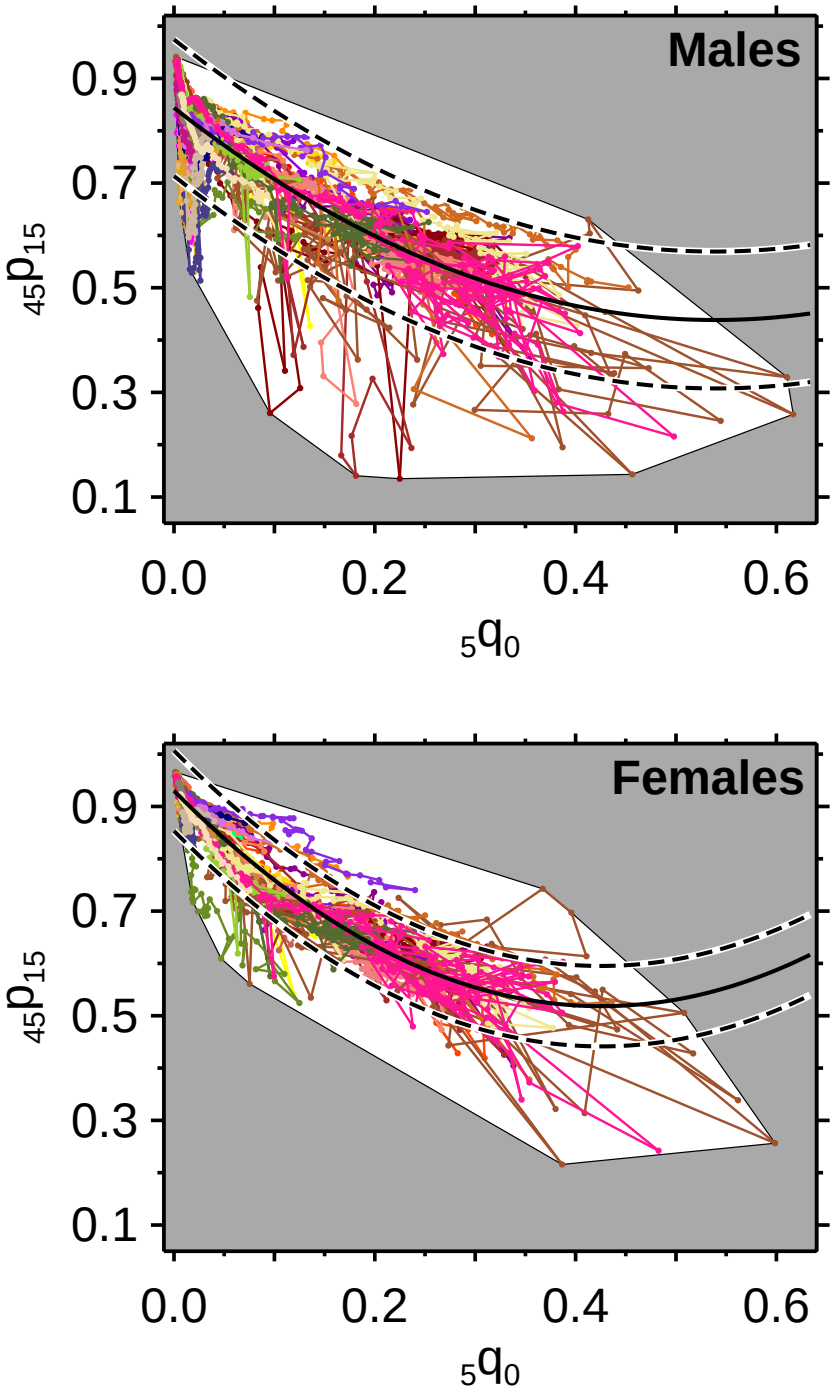


Figure 2: Adult survivorship vs child mortality, by sex. With convex hull, and quadratic regression line (solid) and its associated 95% prediction interval (dashed). Inside the hull, line segments connect chronologically-consecutive points on a per-country basis.

Table 2: **Country peeling**

Peeled country	*N* points contrib. to master CH	*N* points outside peeled hull	Area ratio: peeled hull ÷ master CH	Diameter ratio: PH÷mCH	Number of countries in peeled hull
	(A)	**(B)**	**(C)**	**(D)**	**(E)**
Males					
Belarus	1	5	0.9994	1.0	8
Estonia	1	2	0.9997	1.0	7
Finland	2	2	0.9736	1.0	6
France	1	2	0.9919	1.0	7
Iceland	6	13	0.8060	0.9489	10
Russia	1	4	0.9977	1.0	6
Females					
Belarus	1	3	0.9998	1.0	3
Iceland	8	16	0.6524	0.9324	10
N.Z. (Maori)	4	27	0.9829	1.0	3

Results

i) *Country peeling*: Outlier quantification in the HMD

In using convex hulls to identify outliers, we take a country-centered approach, removing one country at a time and seeing how the resulting hull differs from the *master hull* (i.e., the hull the entire data set, shown in figure 2). This *country peeling* differs from Tukey peeling. In figure 2, the female hull is defined by 13 points from three countries and the male hull by 12 points from six countries. Under Tukey peeling, we would remove these 13 or 12 points, respectively, and examine the modified data set, or, possibly, peel the next hull. With country peeling, we remove entire countries, one at a time, instead of all the vertices of the master convex hull. From the basic information about the master hull, we know *a priori* that the country peeling will have no effect except when the country being peeled is one of those three countries for females (six for males).

When a country is peeled, a new, smaller, hull is calculated to reflect the country-peeled data set. Table 2 lists the component countries of the master hull and the number of points each of those countries contributes to the master hull (column A). The next three columns of table 2 give a

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

counterfactual *as-if-adding* scenario. That is to say, if a country was never in the data set to begin with but was then added, how much of an outlier would it be, relative to the convex hull of the prior mortality relationship. This can be measured by how many points of the country lie outside the prior (i.e., peeled) hull (column B), or as the how large the old hull is, relative to the new one (viz., the as-if-added hull) (column C).

The diameter of a convex hull is the greatest distance between any two vertices. Column D gives the diameter of the peeled hull relative to that of the master hull. This column shows an idiosyncrasy of the mortality relationship in the HMD data, namely that Iceland populates the hull at both ends, high child mortality/low adult survivorship (in the 19th century), and low child mortality/high adult survivorship (in the 21st century). The endpoints of the diameter of the master hull (known as the antipodes) need not be drawn from the same country, but in the HMD data set, they are. Thus, the diameter changes if and only if Iceland is the peeled country; column D highlights this. Column E gives the number of component countries of the peeled hull.

Country peeling is a technique for outlier detection: it answers the question, "if country X is being added to the HMD for the first time, how different would it be from the existing countries?" Table 2 shows that most countries are similar in the mortality relationship. Indeed, for females, only three countries would be flagged as outliers if being added (one at a time) to the HMD for the first time. The most severe outlier, as quantified by table 2, is Iceland. Country peeling for Iceland is shown graphically in figure 3. The inner, peeled, hull (in white, consisting of points from 10 countries for either sex) is smaller in area and has a smaller diameter than the master hull. The take home message of this section is that convex hulls are an effective tool for qualitatively identifying outliers, as well as quantifying their degree of outlierness. Most countries are not outliers, reflecting commonalities in the mortality relationship in the HMD data. The following section looks at intra-country analysis of the mortality relationship by sex, and shows that convex hulls are not just for outlier detection.

ii) Sex differences

Much of population studies concerns time series of demographic phenomena (life expectancy, total fertility rate, etc.). *Gaps and lags* is one way to

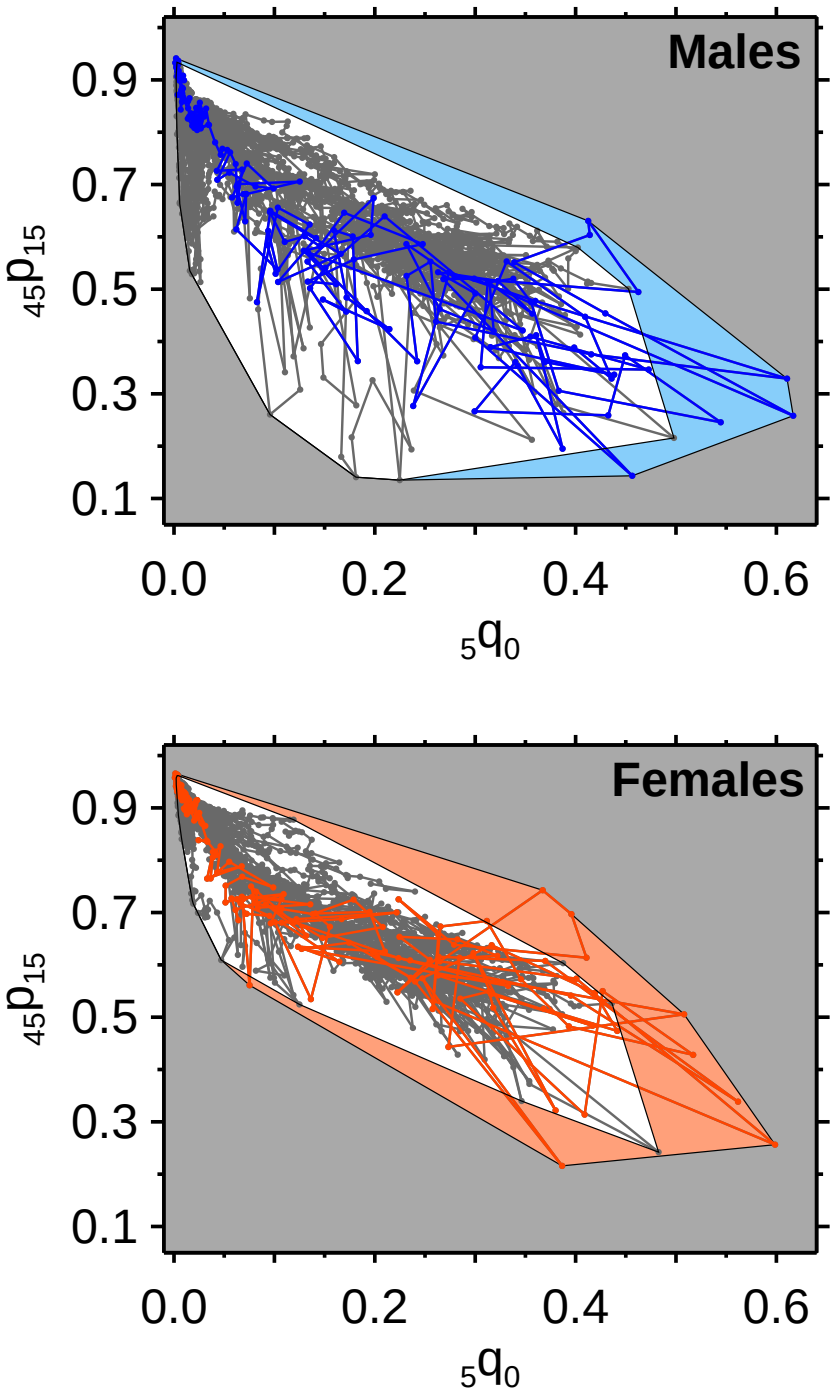


Figure 3: Adult survivorship vs child mortality, by sex. For each sex, two hulls are shown: the outer, colored, hull is the master convex hull (the same as shown in figure 2). The inner, white, hull is the result of peeling Iceland, the data of which is shown in color.

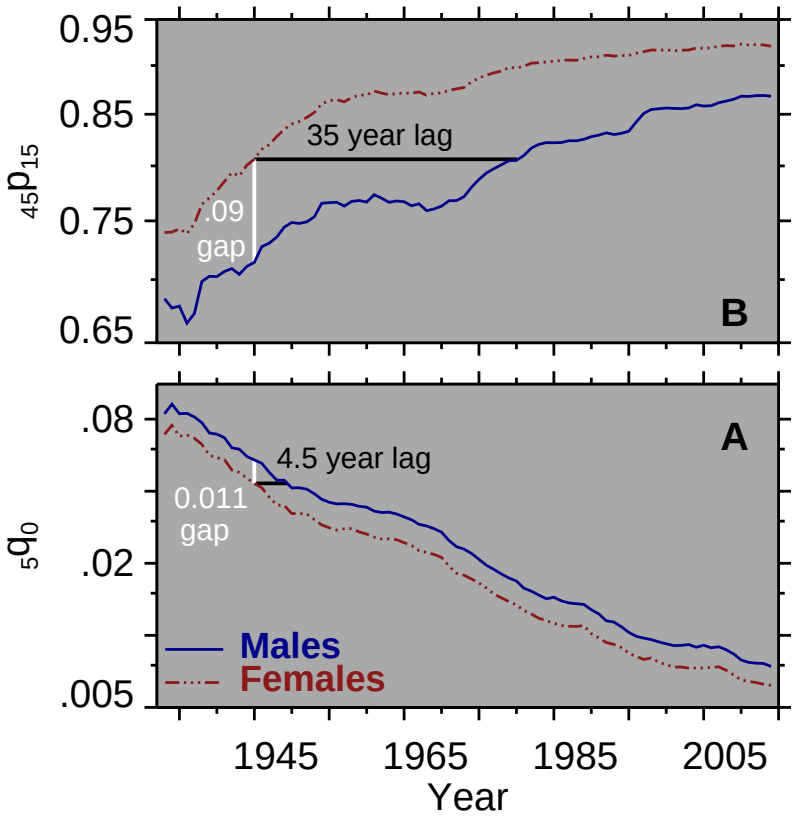


Figure 4: Sex differences in mortality, *gaps and lags* perspective. Female advantage, which is the typical of the overwhelming majority of cases, can be viewed as period gap or as a lag of males, taking them a certain number of years to catch up. A: ${}_5q_0$, child mortality; B: ${}_{45}p_{15}$, adult survivorship. All data from HMD, for the United States.

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

conceptualize the movement of two time series where each one is measuring the same quantity but for different (but related) populations. Figure 4 illustrates this for American males and females; each panel shows a different mortality measure: figure 4A, child mortality (${}_5q_0$); figure 4B, adult survivorship (${}_{45}p_{15}$). The separations between the male and female series can be regarded as a gap (along the vertical axis, shown in white), or as a lag (along the horizontal axis, shown in black). In 1945, the male-female gap in child mortality was 11 per thousand. Or, one could say that the males would take 4.5 more years to achieve the equivalent ${}_5q_0$ as females in 1945 (a lag). Goldstein and Wachter (2006) formalized the gaps and lags framework, using periods and cohorts as the population dichotomy. As we show here without the formalism, this framework can apply also to sex differences.

The basic idea is that any gap can be recast as a lag. In terms of mortality decline, males and females plough the same ground, but the female mortality advantage (or gap) means that males do so later. Figure 5 shows the convex hull approach to this problem. Consider first the "ISL" panel, for Iceland. We see exactly the phenomenon of males following in the path of females: the hulls are substantially overlapping. Of course, given that the time series is the same length for each sex we should not expect total overlap. Males begin the series with mortality levels higher than seen in females, and females end the series with lower mortality than seen in males for the same time interval. Thus, we expect two regions, at opposite ends of the space, where the hulls do not overlap. This is precisely what the graph for Iceland shows.

For the United States (figure 5, "USA"), the convex hull analysis reveals a much different story. Unlike the univariate time series in figure 4, the males' mortality relationship does not follow in the footsteps of the females'. The convex hulls are disjoint, indicating that males and females are not playing a precise game of follow the leader, but are taking different paths. The disparate lags in figure 4 (i.e., 35 years for adult survivorship but only 4.5 years for child mortality) drive the disjointness of the male and female hulls. The long lag for the adult survivorship data is thought to be due principally to behavioral influences, especially tobacco use (see, f.e., Pampel 2002). While a careful read of figure 4 would allow one to predict the divergent paths over time, the hull approach reveals this much more clearly and easily. Note that the high variance in the Iceland data, which goes back to 1838, contributes to the overlap of the hulls. The American data begin in

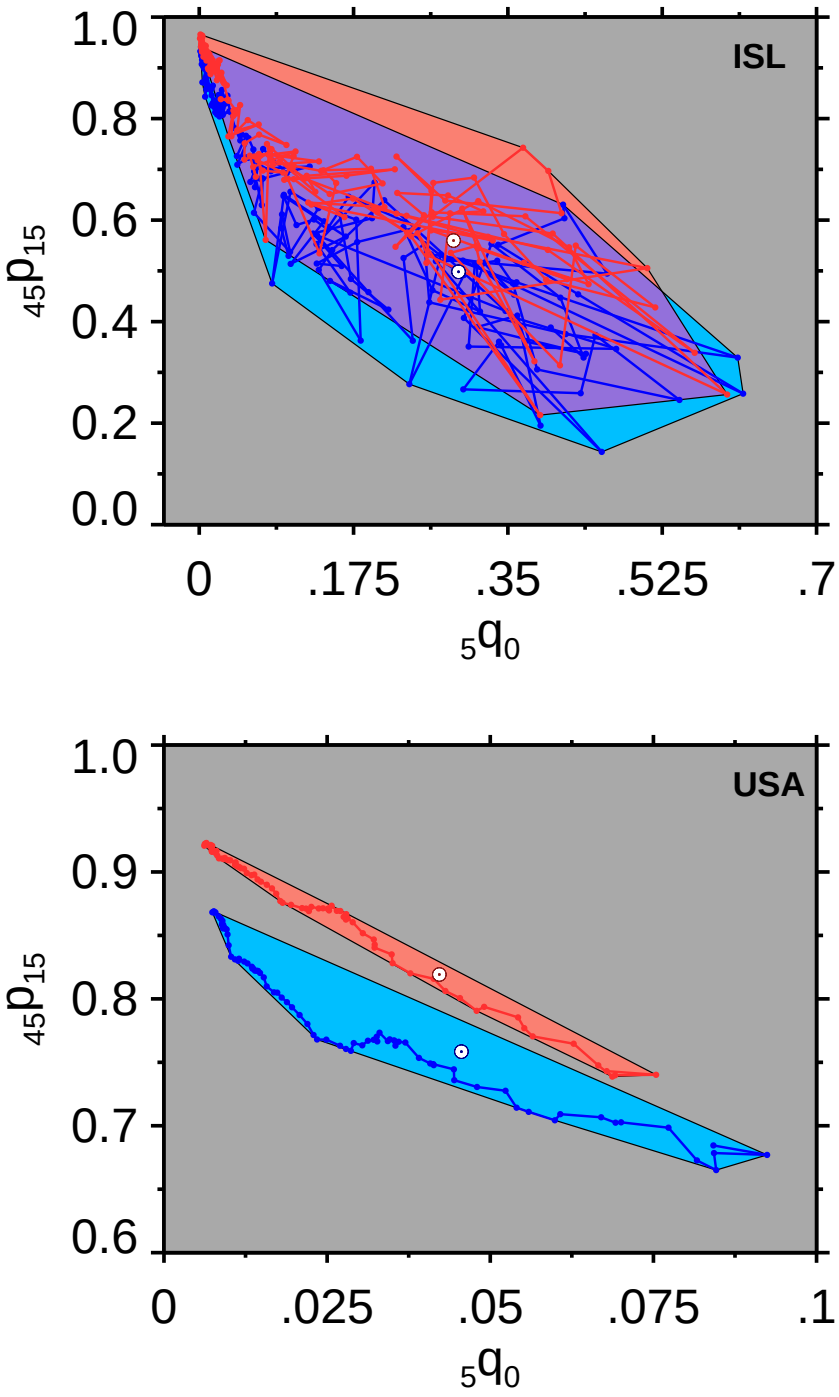


Figure 5: Male (blue) and female (red) sex-specific hulls for the mortality relationship, Iceland (ISL) and United States (USA). Overlap shown in purple. Bullseyes denote the centroids of the hulls.

Table 3: **Male and female hulls: Descriptive and comparative statistics.**

Country	Female to male Area Ratio	Female to male Diameter Ratio	Intersection area as % of M	Intersection area as % of F	% of points in intersection M	% of points in intersection F	Euclidean dist. btwn. centroids	Centroid in opposite hull? M	Centroid in opposite hull? F
	(A)	**(B)**	**(C)**	**(D)**	**(E)**	**(F)**	**(G)**	**(H)**	**(I)**
Australia	0.524	0.874	18.4	35.0	0.0	53.8	0.042	No	No
Austria	0.346	0.777	*disjoint*				0.072	—	
Belarus	0.231	0.295	*disjoint*				0.162	—	
Belgium	0.666	0.985	58.4	87.6	0.0	35.7	0.022	Yes	Yes
Bulgaria	0.292	1.049	0.3	1.1	11.1	0.0	0.074	No	No
Canada	0.439	1.044	34.0	77.4	0.0	31.2	0.016	Yes	Yes
Chile	0.274	0.585	*disjoint*				0.074	—	
Czech Republic	0.337	0.831	*disjoint*				0.086	—	
Denmark	0.936	0.941	78.9	84.3	18.2	40.0	0.039	Yes	Yes
Estonia	0.291	0.368	*disjoint*				0.163	—	
Finland	0.282	0.577	19.8	70.2	0.0	25.0	0.212	No	Yes
France	0.412	0.793	36.4	88.3	0.0	30.0	0.161	No	Yes
E Germany	0.385	0.913	*disjoint*				0.077	—	
W Germany	0.346	0.635	*disjoint*				0.067	—	
Greece	0.308	0.819	*disjoint*				0.059	—	
Hungary	0.286	0.854	*disjoint*				0.109	—	
Ireland	0.484	0.976	12.1	25.0	8.3	27.3	0.039	No	No
Iceland	0.914	1.009	80.8	88.4	20.0	50.0	0.062	Yes	Yes
Israel	0.568	0.613	*disjoint*				0.043	—	
Italy	0.428	0.851	39.6	92.5	0.0	10.0	0.083	No	Yes
Japan	0.809	0.881	19.0	23.5	0.0	37.5	0.044	No	No
Latvia	0.233	0.338	*disjoint*				0.187	—	
Lithuania	0.194	0.336	*disjoint*				0.163	—	
Luxemburg	0.366	0.573	3.9	10.7	0.0	20.0	0.083	No	No
Netherlands	0.645	0.940	61.2	94.9	0.0	72.7	0.032	Yes	Yes
N.Z. (Maori)	1.022	1.472	73.2	71.6	33.3	40.0	0.009	Yes	Yes
N.Z. (Non-Maori)	0.580	0.754	52.5	90.5	0.0	57.1	0.062	Yes	Yes
Norway	0.804	0.905	68.8	85.5	0.0	42.9	0.042	Yes	Yes
Poland	0.220	0.828	*disjoint*				0.118	—	
Portugal	0.635	0.921	*disjoint*				0.086	—	
Russia	0.274	0.307	*disjoint*				0.219	—	
Slovakia	0.368	1.093	*disjoint*				0.097	—	
Slovenia	0.322	0.404	*disjoint*				0.108	—	
Spain	0.555	0.987	42.7	77.0	0.0	9.1	0.049	Yes	Yes
Sweden	0.883	0.988	83.3	94.3	0.0	50.0	0.035	Yes	Yes
Switzerland	0.632	0.890	56.8	89.8	0.0	33.3	0.066	Yes	Yes
Taiwan	0.907	0.969	*disjoint*				0.079	—	
England and Wales	0.296	0.815	25.6	86.4	0.0	57.1	0.123	No	Yes
Scotland	0.748	0.959	44.1	59.0	8.3	50.0	0.042	Yes	Yes
Northern Ireland	0.661	1.131	52.0	78.7	0.0	40.0	0.007	Yes	Yes
Ukraine	0.225	0.230	*disjoint*				0.166	—	
United States	0.332	0.896	*disjoint*				0.061	—	

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

1933 and show less variance, giving thinner hulls. Of the twenty countries with disjoint hulls, (table 3), all have data beginning after the Second World War, except Portugal (1940) and the United States (1933) (cf. table 1). With improvements in nutrition, the advent of antibiotic drugs, and so on, the postwar mortality regime is lower variance (at least in the HMD member countries), which tends to favor these disjoint hulls. However, note that Austria (1947), Bulgaria (1947), Ireland (1950), Japan (1947), Luxembourg (1960), and New Zealand/Maori (1948) are all exclusively postwar data, yet have overlapping hulls, and therefore the overlap is not exclusively driven by noisy prewar data.

Table 3 compares the male and female hulls on a per-country basis for all the countries. Statistics provided are area (column A) and diameter (column B) ratios, the intersection area as a percentage of each single-sex hull (columns C,D), the percentage of points of each hull within their intersection (columns E,F), the distance between the centroids of each hull (column G), and indicators for whether the hull centroids lie inside the opposite-sex hull (columns H,I). Almost half the hulls (20/42) are disjoint across the sexes, indicating that male mortality decline, as measured by the mortality relationship, does not follow in the footsteps of the female mortality decline.

The take home message of this section is that when mortality is summarized in more than one dimension, a male-female gap cannot be construed to be the same as a lag. The convex hull of the male mortality relationship traced out over time is in many cases disjoint with the same hull for females from the same national population.

iii) Period and cohort

This section presents a similar analysis as that immediately above, but we examine periods and cohorts, by country and sex. Finland is shown in figure 6 as an example, and the appendix (p. 20) gives graphs for all the other HMD countries with cohort data. The Finnish data illustrate a number of interesting traits. The male period hull is much bigger than the male cohort hull, and envelops it.

Table 4 gives some descriptive statistics for the cohort and period hulls. As before, we give ratios of one hull to another (in this case, period to cohort). In addition to area (column A) and diameter (column B) we include

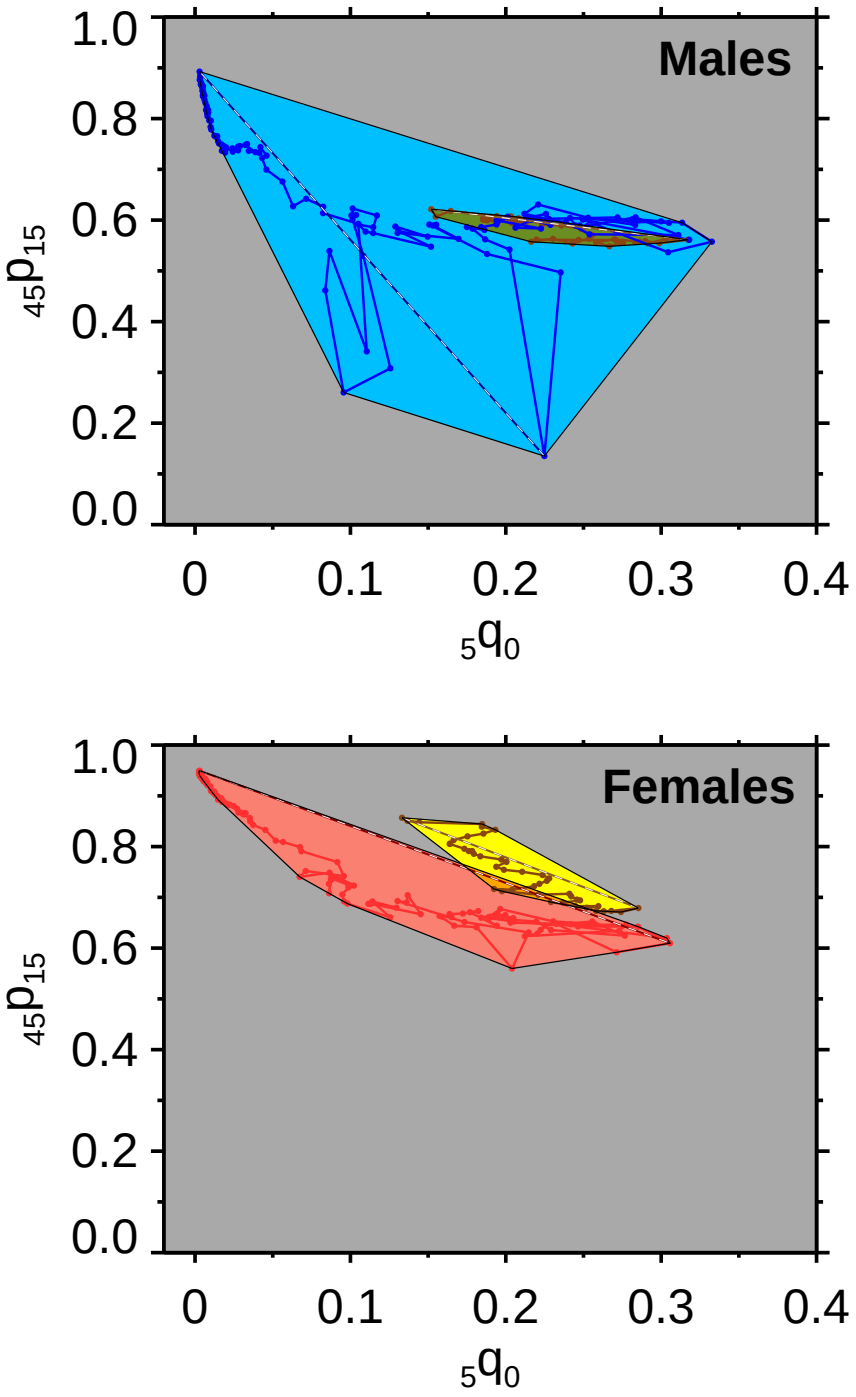


Figure 6: Male (blue) and female (red) convex hulls for the period mortality relationship for Finland. The cohort mortality relationship is shown in yellow (or green/orange or where it overlaps with male/female [respectively] period data). Dashed lines denote the diagonals of the hulls.

perimeter (column C) as a linear measure of hull size. Column D gives the percentage of the cohort hull that lies outside the period hull. Most have some overlap, but four hulls (out of 22) are disjoint, including both sexes for the Swiss data. In only one case (Finnish males) is the cohort hull enveloped by the period hull (although for French and Swedish males, 1% or less of the area of the cohort hull lies outside the period hull). We also introduce the *diagonal correlation* (column E), or the cosine of the angle between the diagonals. The diagonal correlation provides a rough but useful dimensionless measure of how parallel, so to say, the two hulls are. Of course, while diagonals may be perfectly parallel — and five of the hulls indeed have perfect diagonal correlation (modulo rounding up to 1.0) — the hulls themselves are polygons which do not have correlations in the traditional sense. Nonetheless, the Finnish example demonstrates the utility of the diagonal correlation measure. For males, the diagonal correlation is 0.59 while for females it is 1.0, after rounding. Compare these to the hulls illustrated in figure 6, where the alignment (using a non-rigorous, intuitive meaning) of the two hulls seems much better for females.

Similar to the analysis of the male and female mortality relationship, the convex hull analysis of periods and cohorts shows that when mortality is summarized as the mortality relationship, period data does not follow cohort data with a lag.

Conclusion

The goal of this work is to introduce convex hull analysis to demography, mostly as tool for exploratory data analysis (Tukey, 1977). Further refinements are possible, such as hypothesis testing (see, f.e., Rogers 1978 — although there are difficulties using convex hulls as inferential tools when the quantities being analyzed don't follow standard distributions such as normal, which is often the case in demography).

This is the first work that we are aware that makes extensive use of convex hulls as an analytic tool or framework for population statistics. We believe the foregoing analysis demonstrates that convex hull analysis can sometimes be a useful tool for demographers and we encourage population scientists to consider their use.

Table 4: **Cohort and period hulls: Descriptive and comparative statistics.**

Country	sex	period to cohort ratio			% of coh.	diagonal
		area	diameter	perimeter	outside per.	correlation
		(A)	**(B)**	**(C)**	**(D)**	**(E)**
Denmark	M	2.341	1.869	1.741	*disjoint*	0.980
	F	2.707	1.677	1.608	93.5	1.000
Finland	M	26.725	4.487	4.624	*enveloped*	0.592
	F	3.600	1.944	2.002	88.4	1.000
France	M	5.212	2.413	2.371	1.0	0.988
	F	3.438	1.808	1.799	68.5	1.000
Iceland	M	2.913	1.364	1.492	12.3	0.990
	F	3.101	1.530	1.598	19.5	0.962
Italy	M	6.451	3.591	2.978	56.0	0.773
	F	4.353	2.464	2.470	*disjoint*	0.985
Netherlands	M	3.156	2.120	1.985	98.8	0.956
	F	2.150	1.812	1.672	97.5	0.989
Norway	M	4.290	1.775	1.892	7.4	0.986
	F	2.998	1.438	1.566	34.0	1.000
Sweden	M	4.386	1.766	1.760	0.5	0.999
	F	5.126	1.799	1.775	3.5	0.999
Switzerland	M	6.815	1.919	2.152	*disjoint*	0.947
	F	4.890	1.718	1.858	*disjoint*	0.973
England and Wales	M	4.616	2.131	2.118	76.3	0.944
	F	1.296	1.587	1.504	95.3	1.000
Scotland	M	3.190	2.091	1.977	88.6	0.992
	F	2.082	1.670	1.622	83.0	0.992

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

Works Cited

Aldous, David J., Bert Fristedt, Philip S. Griffin, and William E. Pruitt. 1991. "The number of extreme points in the convex hull of a random sample." *Journal of Applied Probability* 28(2):287–304.

Barbieri, Magali, John R. Wilmoth, Vladimir M Shkolnikov, Dana Glei, Domantas Jasilionis, Dmitri Jdanov, Carl Boe, Timothy Riffe, Pavel Grigoriev, and Celeste Winant. 2015. "Data resource profile: The Human Mortality Database (HMD)." *International Journal of Epidemiology* 44(5):1549–1556.

Barnett, V. 1976. "The ordering of multivariate data." *Journal of the Royal Statistical Society, Series A* 139(3):318–355.

Bebbington, A. C. 1978. "A method of bivariate trimming for robust estimation of the correlation coefficient." *Applied Statistics* 27(3):221–226.

Blackwell, Paul. 1992. "Convex hulls of selected subsets of a Poisson process." *Journal of Applied Probability* 29(4):814–824.

Cornwell, William K., Dylan W. Schwilk, and David D. Ackerly. 2006. "A trait-based test for habitat filtering: Convex hull volume." *Ecology* 87(6):1465–1471.

de Berg, Mark, Otfried Cheong, Marc van Kreveld, and Mark Overmars. 2008. *Computational geometry: Algorithms and applications*. Springer, Berlin, third ed.

DeGroot, Morris H. and Mark J. Schervish. 2002. *Probability and statistics*. Addison-Wesley, Reading, MA, third ed.

Eddy, William F. 1980. "The distribution of the convex hull of a Gaussian sample." *Journal of Applied Probability* 17(3):686–695.

Efron, Bradley. 1965. "The convex hull of a random set of points." *Biometrika* 52(3/4):331–343.

Galloway, Patrick R. 1994. "A reconstruction of the population of North Italy from 1650 to 1881 using annual inverse projection with comparisons to England, France, and Sweden." *European Journal of Population* 10(3):223–274.

Georgakis, Giannis and George Tziafetas. 1982. "Stochastic projection of the population distribution in Greece." *European Demographic Information Bulletin* 13(3):120–134.

Getz, Wayne M. and Christopher C. Wilmers. 2004. "A local nearest-neighbor convex-hull construction of home ranges and utilization distributions." *Ecography* 27(4):489–505.

Goldstein, Joshua R. and Kenneth W. Wachter. 2006. "Relationships between period and cohort life expectancy: Gaps and lags." *Population Studies* 60(3):257–269.

NBER Cohort Studies 2017 meeting submission. Pls contact authors for latest ver.

Goldstone, J. A. 1986. "The demographic revolution in England: A re-examination." *Population Studies* 40(1):5–33.

Hodge, Victoria and Jim Austin. 2004. "A survey of outlier detection methodologies." *Artificial Intelligence Review* 22(2):85–126.

Huber, Peter J. 1972. "Robust statistics: A review." *Annals of Mathematical Statistics* 43(4):1041–1067.

Hueter, Irene. 1994. "The convex hull of a normal sample." *Advances in Applied Probability* 26(4):855–875.

Human Mortality Database. 2016. http://www.mortality.org/. Accessed 15 August 2016.

Massé, Bruno. 2000. "On the LLN for the number of vertices of a random convex hull." *Advances in Applied Probability* 32(3):675–681.

Pampel, Fred C. 2002. "Cigarette use and the narrowing sex differential in mortality." *Population and Development Review* 28(1):77–104.

Preparata, Franco P. and Michael Ian Shamos. 1985. *Computational geometry: An introduction*. Springer, New York.

Rogers, L. C. G. 1978. "The probability that two samples in the plane will have disjoint convex hulls." *Journal of Applied Probability* 15(4):790–802.

Snedecor, George W. and William G. Cochran. 1989. *Statistical methods*. Iowa State University Press, Ames, eighth ed.

Snyder, Timothy Law and J. Michael Steele. 1993. "Convex hulls of random walks." *Proceedings of the American Mathematical Society* 117(4):1165–1173.

Suri, Subhash, Kevin Verbeek, and Hakan Yıldız. 2013. "On the most likely convex hull of uncertain points." In Hans L. Bodlaender and Giuseppe F. Italiano (eds.), *Algorithms – ESA 2013: 21st annual European symposium, Sophia Antipolis, France, September 2–4, 2013. Proceedings*, pp. 791–802. No. 8125 in Lecture Notes in Computer Science, Springer, Berlin.

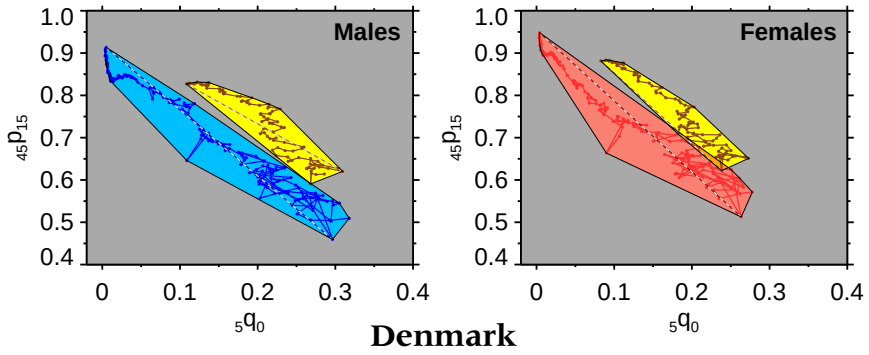
Tukey, John W. 1975. "Mathematics and the picturing of data." In Ralph D. James (ed.), *Proceedings of the International Congress of Mathematicians, Vancouver 1974*, vol. 2, pp. 523–531.

———. 1977. *Exploratory data analysis*. Addison-Wesley, Reading, MA.

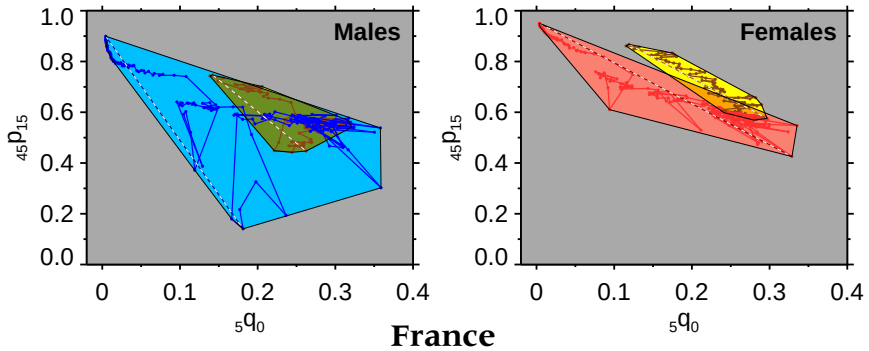
Wrigley, E. A. and R. S. Schofield. 1981. *The population history of England, 1541–1871: A reconstruction*. Edward Arnold, London.

Appendix: Period-cohort hulls for all countries

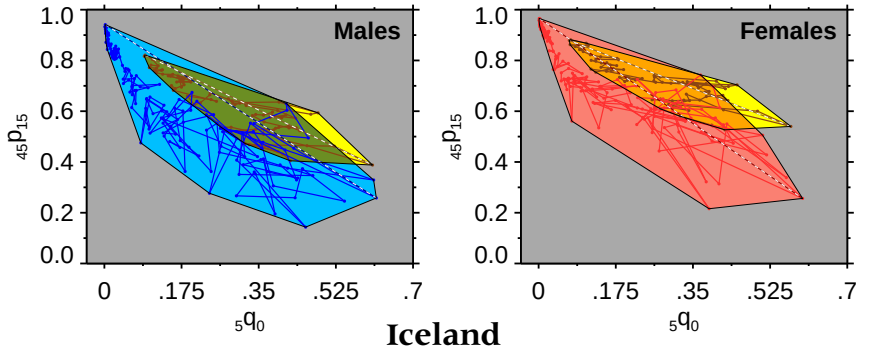
This appendix shows the period-cohort hulls for all countries apart from Finland (which is shown in figure 6 in the main text).



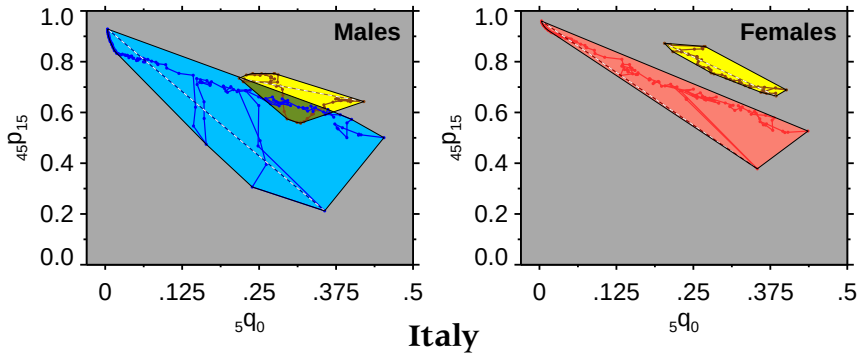
Denmark



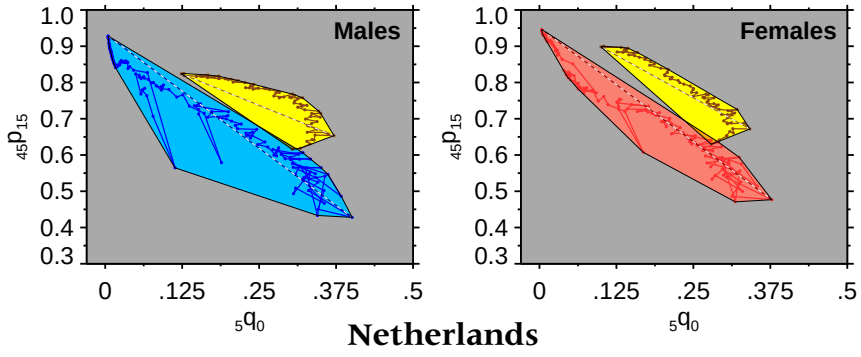
France



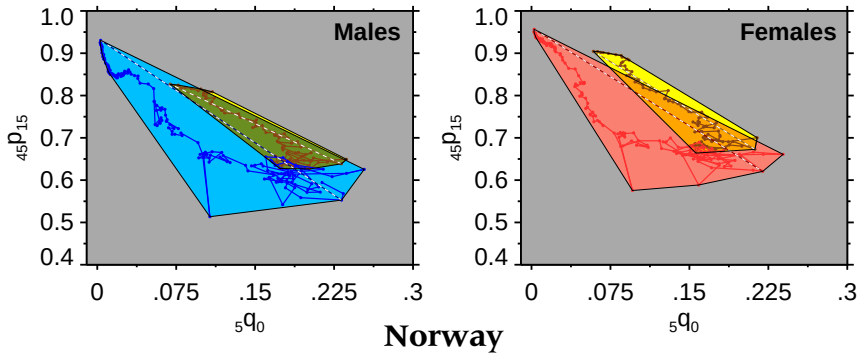
Iceland



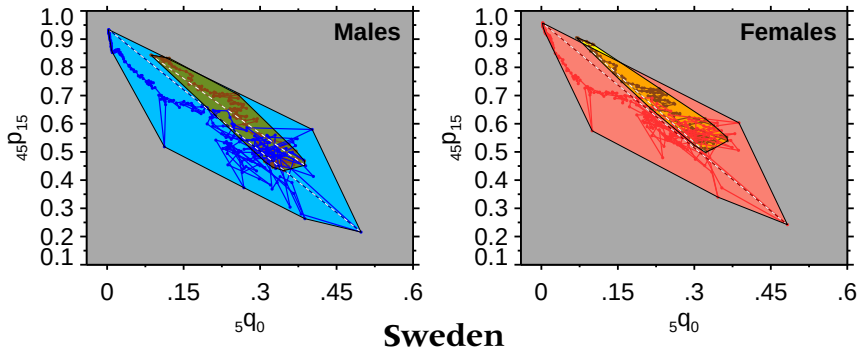
Italy



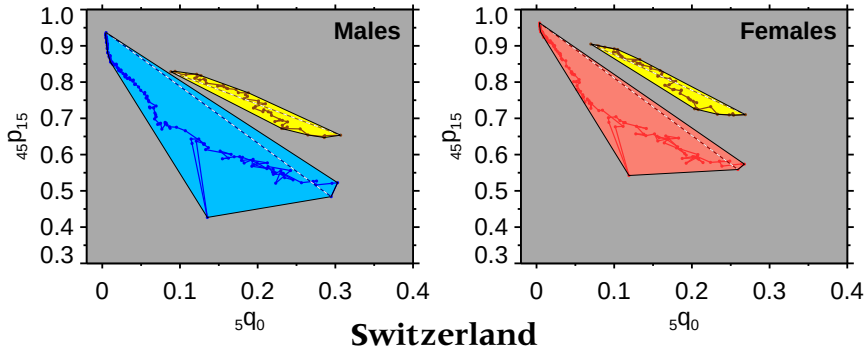
Netherlands



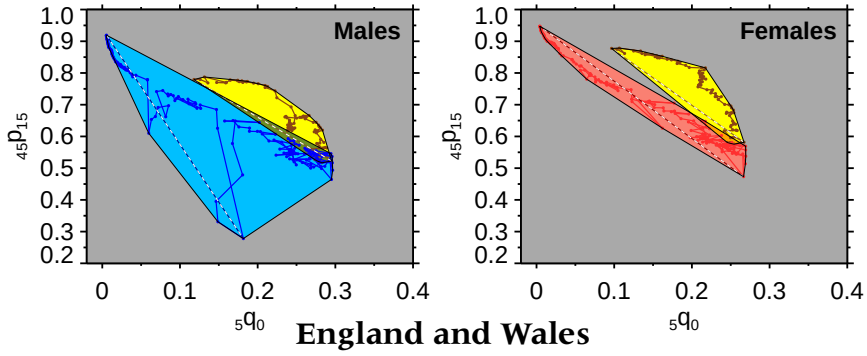
Norway



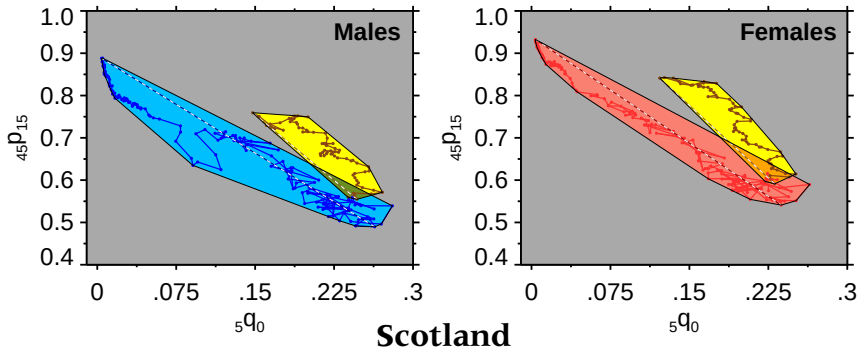
Sweden



Switzerland



England and Wales



Scotland