Journal of *Xenobiotics*

MDPI

Article

Joint Associations of Per- and Polyfluoroalkyl Substance Mixtures with All-Cause, Cardiovascular and Cancer Mortality

Elizabeth Olarewaju [1,2] and Emmanuel Obeng-Gyasi [1,2,*]

[1] Department of Built Environment, North Carolina A&T State University, Greensboro, NC 27411, USA
[2] Environmental Health and Disease Laboratory, North Carolina A&T State University, Greensboro, NC 27411, USA
* Correspondence: eobenggyasi@ncat.edu; Tel.: +1-336-285-3132

Abstract

Per- and polyfluoroalkyl substances (PFAS) occur as correlated mixtures, yet their joint association with mortality remains poorly characterized. We followed 5964 adults from the National Health and Nutrition Examination Survey 2003 to 2010 over 69,256 person-years to 31 December 2019, during which 1119 died, 344 from cardiovascular disease and 235 from malignant neoplasms. Joint associations were estimated by quantile g-computation within Cox models, supported by competing-risks, spline and Bayesian kernel machine models. A simultaneous one-quartile increase in all six analytes gave no clear evidence of association with cardiovascular mortality (hazard ratio 0.92, 95 percent confidence interval 0.79 to 1.07) or cancer mortality (0.90, 0.75 to 1.08). All-cause mortality was inversely associated (0.83, 0.77 to 0.90), and survey weighting left the estimate unchanged (0.82, 0.74 to 0.90). The joint exposure-response was nonlinear ($p = 0.009$) but additive across analytes ($p = 0.299$), and the signal lay at the lower end of the exposure range, where the hazard ratio reached 1.49 (1.29, 1.72) with all six analytes at their joint 10th percentile relative to their 50th. The inverse estimate was confined to the first decade of follow-up and attenuated most on adjustment for serum albumin and glomerular filtration rate (0.87, 0.80 to 0.94). Smoking, household income and seafood consumption did not account for it. These data give no clear evidence that PFAS mixtures at general-population concentrations raise cardiovascular or cancer mortality, although power for the cause-specific endpoints was limited.

Keywords: PFAS; PFOA; PFOS; environmental mixtures; all-cause mortality; cardiovascular mortality; cancer mortality; NHANES; survival analysis; Bayesian kernel machine regression; quantile g-computation

Check for updates

Academic Editor: Lili Tang

Received: 26 October 2026
Revised: 15 November 2026
Accepted: 22 November 2026
Published: 25 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Per- and polyfluoroalkyl substances (PFAS) are synthetic fluorinated organic compounds manufactured since the 1940s and used in non-stick cookware, stain-resistant fabrics, firefighting foams and food packaging [1,2]. Strong carbon-fluorine bonds make them resistant to heat, water and oil, and the same stability has made them ubiquitous contaminants of air, water, soil, wildlife and human populations [3,4]. Long-chain members bind strongly to serum albumin- and liver fatty acid-binding proteins, which slows elimination [5]. Estimated human serum half-lives are approximately 2.7 to 3.5 years for perfluorooctanoic acid (PFOA), 3.4 to 5.7 years for perfluorooctane sulfonate (PFOS) and 5.3 years for perfluorohexane sulfonic acid (PFHxS) [6,7], so body burden persists for years after exposure ceases. PFOA, PFOS, perfluorononanoic acid (PFNA), PFHxS, perfluorodecanoic acid (PFDeA) and perfluoroundecanoic acid (PFUnDA) are the most extensively

studied members of the class [8,9]. Production of PFOA and PFOS was phased out in the United States during the 2000s and formalized under the Stockholm Convention on Persistent Organic Pollutants in 2009 and 2019, respectively [1,10], but legacy contamination persists across all environmental matrices [1,3].

Human exposure occurs through contaminated drinking water, particularly near manufacturing facilities, military installations using aqueous film-forming foam and industrial discharge sites [11,12]; through diet, especially fish, shellfish and food in treated packaging [2,13,14]; and through inhalation, dermal contact and indoor dust [13]. Biomonitoring in the National Health and Nutrition Examination Survey (NHANES) detected PFOA, PFOS, PFHxS and PFNA in more than 98 percent of United States adults during 2003 to 2014, confirming that co-exposure to PFAS (per- and polyfluoroalkyl substances) mixtures, rather than to any single compound, is the rule in the general population [8,9].

PFAS bind the peroxisome proliferator-activated receptors, the pregnane X and constitutive androstane receptors and thyroid hormone receptors, dysregulating lipid metabolism, glucose homeostasis, inflammatory signaling and endocrine function [15–17], and they compete with endogenous fatty acid ligands on albumin- and liver fatty acid-binding proteins [5]. Human studies report dyslipidaemia, liver dysfunction, thyroid disorders and attenuated antibody responses to vaccination [18–20], and serum concentrations have been associated with insulin resistance, type 2 diabetes, hypertension and metabolic syndrome [21,22]. Oxidative stress and systemic inflammation, mechanisms central to atherosclerosis and to multiple malignancies, have been proposed as the shared pathway [23–25]. In its 2023 evaluation the International Agency for Research on Cancer (IARC) classified PFOA as carcinogenic to humans (Group 1) and PFOS as possibly carcinogenic (Group 2B) [26,27].

The epidemiological evidence on mortality is divided. A population-based study of the Veneto Region of Italy, where about 150,000 people were exposed through drinking water, found excess all-cause and cardiovascular mortality across nearly four decades of follow-up [28], and a cohort from a contaminated Swedish municipality observed elevated cardiovascular morbidity and mortality [29]. Two large Scandinavian prospective cohorts, in contrast, found no overall association between plasma PFAS and incident cardiovascular disease, with an inverse association for PFOS in men [30]. For cancer, evidence is most consistent for kidney and testicular sites in occupational cohorts and in communities exposed through drinking water [31–33], while signals for breast, thyroid, prostate, bladder and colorectal cancer are less consistent and mixture effects on cancer mortality have not been rigorously examined [31,34]. The limited cohort literature in representative general-population samples suggests elevated all-cause mortality with higher body burdens, but single-compound frameworks and inadequate control for the correlation among analytes preclude firm conclusions [29,33].

Two features of this evidence motivate the present analysis. First, PFAS occur together and are strongly correlated, so single-compound estimates cannot be separated from the mixture in which they sit, and methods designed for correlated mixtures are required [35]. Second, serum concentrations changed substantially across the NHANES (National Health and Nutrition Examination Survey) cycles carrying PFAS measurements, and because those cycles also differ in how long participants have been followed, calendar period is a common cause of exposure and of the number of deaths observed, so any pooled-cycle analysis must handle this explicitly. The objectives were to estimate the joint association of the six-analyte serum PFAS mixture with all-cause, cardiovascular and cancer mortality using time-to-event models with individual follow-up time; to test whether any association reflects reverse causation, through lag exclusions and period-specific estimates; to characterize the exposure-response surface and identify which analytes carry posterior selection support; and to establish whether the findings are robust to competing risks, to the

complex survey design and to adjustment for smoking, socioeconomic position, lifestyle and baseline health.

2. Materials and Methods

2.1. Study Population and Mortality Follow-Up

Data came from the National Health and Nutrition Examination Survey (NHANES) 2003–2004, 2005–2006, 2007–2008 and 2009–2010, the four cycles in which the six-analyte per- and polyfluoroalkyl substance (PFAS) panel analyzed here was measured in the same subsample [36]. Serum PFAS are measured in a one-third subsample of examined participants, which determines the size of the analytic sample. Participants were eligible if they were aged 20 years or over at examination, were eligible for mortality linkage, and had complete serum concentrations for all six analytes. Of 7060 linkage-eligible participants with complete PFAS, 6535 were aged 20 or over and 6431 also had body mass index (BMI) and educational attainment recorded. Primary models were fitted in the 5964 of these participants with complete data for every covariate in the primary model; the family poverty–income ratio (PIR) was the limiting variable. Estimates in the full sample of 6431 under the covariate set that excludes smoking and the PIR (poverty–income ratio) are reported alongside them.

Mortality was ascertained from the National Center for Health Statistics (NCHS) public-use linked mortality files, which match NHANES (National Health and Nutrition Examination Survey) records to the National Death Index through 31 December 2019 [37]. Vital status and person-months of follow-up from examination were recorded for every mortality-eligible participant and were complete for all participants analyzed. In the primary sample, follow-up totalled 69,256 person-years, with a median of 141 months (interquartile range 119 to 170, range 1 to 204). Public-use follow-up times are subject to confidentiality-protective perturbation for a small number of records, which adds noise but does not bias the comparisons reported here. All-cause mortality was defined as death from any cause. Cardiovascular mortality was defined from the underlying cause of death as diseases of the heart or cerebrovascular disease, and cancer mortality as malignant neoplasms. Analyses using person-months to interview in place of person-months to examination are reported as a sensitivity analysis.

2.2. Exposure Measurement

Serum PFOS (perfluorooctane sulfonate), PFOA (perfluorooctanoic acid), PFNA (perfluorononanoic acid), PFHxS (perfluorohexane sulfonic acid), PFDeA (perfluorodecanoic acid) and PFUnDA (perfluoroundecanoic acid) were measured by online solid-phase extraction with high-performance liquid chromatography and tandem mass spectrometry and reported in ng/mL [36]. Concentrations below the limit of detection (LOD) are reported in the public files as the LOD (limit of detection) for that analyte and cycle divided by the square root of two, and those substituted values were analyzed as supplied, with no further imputation and no exclusion of participants on the basis of detection status. The limits were 0.2 ng/mL for the 2003–2006 cycles and 0.1 ng/mL for 2007–2010, with 0.082 ng/mL for PFNA in one cycle. Detection frequencies in the analytic sample differed markedly across analytes: 99.8 percent for PFOS, 99.6 percent for PFOA, 99.0 percent for PFNA and 98.2 percent for PFHxS, against 70.1 percent for PFDeA and 60.2 percent for PFUnDA (Table 2). Because the two lowest-detection analytes carry a large block of tied substituted values, the mixture was also restricted to the four analytes with near-complete detection as a sensitivity analysis, and the quartile analyses were repeated with cut points assigned so that tied values fall in the same quartile.

Concentrations were transformed to base-two logarithms for the per-doubling analyses and scored into quartiles for the mixture analyses. Quartile cut points were computed once, across the whole analytic sample, and not within survey cycles, within covariate strata or within the complete-case subset, so that a quartile score has the same meaning for every participant. Each participant contributed a single serum measurement, taken at examination, which indexes recent and medium-term exposure rather than exposure over the whole follow-up period; the implications are examined in the lag and period-specific analyses and in Section 4.

2.3. Covariates

Primary models were adjusted for age at examination in years, sex, race and Hispanic origin, educational attainment, body mass index (BMI), survey cycle, smoking status and the family poverty–income ratio (PIR). Smoking status was classified as never, former or current from self-reported lifetime and current cigarette use. The PIR (poverty–income ratio) is the ratio of family income to the federal poverty threshold and was used in preference to the household income bracket because it is defined in the same way in all four cycles. Availability of every covariate by survey cycle is given in Table 3. An extended model additionally adjusted for alcohol use, defined as at least 12 drinks in any one year, any moderate or vigorous leisure-time physical activity, and self-reported cardiovascular disease, cancer, diabetes and hypertension at examination. Serum cotinine replaced self-reported smoking in a sensitivity analysis. Serum albumin and the estimated glomerular filtration rate (eGFR), calculated with the 2021 Chronic Kidney Disease Epidemiology Collaboration creatinine equation, were added as markers of baseline physiological status in a further sensitivity analysis, because PFAS (per- and polyfluoroalkyl substances) bind albumin and are cleared by the kidney. The household income bracket was used for the socioeconomic gradient analysis. Dietary intake was characterized from the dietary interview as the number of fish and shellfish meals in the past 30 days and total energy intake, which were used in a further sensitivity analysis because seafood consumption is the dietary pathway most likely to carry both exposure and socioeconomic position. Analyses were restricted to participants with complete data for the covariates in each model.

2.4. Statistical Analysis

The primary analysis was a Cox proportional hazards model with time since examination as the time scale, using individual follow-up months. The joint association of the mixture was estimated by quantile g-computation within that model [38]. Each analyte was scored into sample quartiles, all six quartile scores entered a single model with the covariates, and the joint estimate is the sum of the six coefficients, that is, the effect of moving every analyte up by one quartile at once. Its variance was obtained from the full covariance matrix of those six coefficients, so it accounts for the correlation among them, and analyte weights are each coefficient expressed as a share of the positive or negative partial sum. Quantile g-computation was used in preference to weighted quantile sum regression because it imposes no directional constraint, so a single model answers the question and there is no need to fit and then select between a positive-direction and a negative-direction model [38,39]. Parallel models used base-two logarithms of the concentrations to give per-doubling estimates. Ties were handled by the Breslow approximation, and the proportional hazards assumption was assessed by regressing scaled Schoenfeld residuals on rank-transformed time.

Cause-specific mortality was modeled in two ways: a cause-specific Cox model censoring deaths from other causes, and a Fine–Gray subdistribution hazard model treating them as competing events, with the risk set weighted by the estimated censoring distribution

and confidence intervals from 400 bootstrap resamples of participants. The two answer different questions, the first about the rate among those still alive and the second about the cumulative incidence in the presence of competing death, and both are reported.

The possibility of reverse causation was examined through analyses that the time-to-event data make possible. The primary lag analysis is a landmark analysis. At a landmark of two years the risk set was restricted to participants still alive and under observation at 24 months, follow-up began at the landmark, and time was measured from it, so that neither the person-time nor the deaths occurring before the landmark contribute; the same construction was used at a five-year landmark. This avoids the immortal person-time that would arise from deleting early deaths while retaining the early follow-up of survivors. For comparison, the simpler construction in which early deaths are removed and follow-up is still measured from examination is also reported. Separately, follow-up was split at ten years and the association estimated within each period, with participants censored at ten years in the first period and admitted at ten years in the second. Because the two periods use disjoint risk sets the partial likelihood factorizes across them, so the two estimates are independent and their difference was tested with a Wald statistic. An association arising because illness preceding death lowers serum PFAS should weaken under a landmark and should be concentrated in early follow-up; an association reflecting long-term exposure should not.

Because serum PFAS were measured in a one-third subsample, design-based analyses used the PFAS subsample weight, recorded as WTSA2YR in 2003–2006 and WTSC2YR in 2007–2010, divided by four for the four pooled cycles following the NHANES analytic guidelines [40], with the variance clustered on the primary sampling unit within survey stratum. Both the design-based and the unweighted estimates are reported for every endpoint and both exposure scales in Table 4. The unweighted models are used as the reference specification, for three reasons. The analytic sample is restricted to participants with a complete covariate record, so the subsample weights no longer reproduce the target population and the design-based estimates are not unbiased for a population quantity either. The two sets of estimates agree closely, so the design-based and the unweighted analyses support the same conclusions, and the design-based intervals are the wider of the two. And the competing-risks, spline and kernel machine analyses cannot be weighted on the same basis, so a single unweighted reference specification keeps every analysis in the paper comparable. The conclusions do not depend on the choice.

The shape of the joint exposure-response was examined on the hazard scale, by partial likelihood, so that the conclusions about shape do not depend on any approximation of the survival outcome. Three nested Cox models were fitted: one linear and additive in the six base-two logarithmic concentrations; one with natural cubic splines on three degrees of freedom in each analyte, which tests departure from linearity; and one adding all fifteen pairwise products, which tests departure from additivity. The models were compared by likelihood ratio tests, and the joint exposure-response was traced by moving all six analytes together from their joint 10th to their joint 90th percentile under the spline model, referenced to all six at their joint 50th percentile.

Bayesian kernel machine regression (BKMR) was used to explore the exposure-response surface [41,42]. The specification used here models martingale residuals from a Cox model fitted to the primary covariates alone as a Gaussian outcome, so that the exposure surface is estimated against the part of the mortality experience the covariates do not explain and individual follow-up time enters through the residual; a probit specification on the binary outcome is reported alongside for comparability with the earlier mixture literature. This is an exploratory residual-based analysis and not a survival kernel machine model. Martingale residuals are not Gaussian: they are bounded above by one, unbounded

below, and skewed, with skewness of −0.65 for all-cause, +2.28 for cardiovascular and +3.96 for cancer mortality in these data, and their variance depends on the underlying survival process. The Gaussian likelihood is therefore a working approximation, and the credible intervals and posterior inclusion probabilities (PIPs) it produces are descriptive summaries of the fitted surface rather than calibrated inferential statements. The Cox models are the inferential analysis, the spline models carry the inference about shape, and the kernel machine results show the surface and whether it is consistent with them.

Both kernel machine specifications used a Gaussian kernel with component-wise variable selection across all six analytes and a 50-knot Gaussian predictive process approximation, with knots chosen by k-means on the standardized base-two logarithmic exposures, so that the full cohort could be used rather than a computational subsample. Priors were an inverse-gamma (0.001, 0.001) on the residual variance, a gamma (10, rate 0.1) on the ratio of the kernel variance to the residual variance, a gamma (1, rate 2) truncated at 3 on the analyte-specific kernel coefficients, with half-normal random-walk proposals whose scale was adapted during burn-in, and a Bernoulli (0.5) prior inclusion probability for each analyte. The truncation keeps the kernel in the smooth regime and prevents the degenerate solution in which the kernel matrix approaches the identity and the surface absorbs independent noise. Three chains were run per model from seeds 101, 202 and 303. The residual-scale models were run for 4000 iterations with the first 1500 discarded and the remainder thinned by two, with the selection indicators held fixed for the first 400 iterations so that the kernel coefficients and the variance parameters could adapt before any component was dropped. The probit models carry less information per observation, because a binary outcome replaces the follow-up time, and were run for 9000 iterations with the first 4000 discarded, thinned by four, with the selection indicators held fixed for the first 1500 iterations. Convergence was assessed by the split Gelman–Rubin statistic and the effective sample size for the overall mixture contrast, the kernel variance ratio and each analyte-specific kernel coefficient.

Reported quantities are the PIP (posterior inclusion probability) for each analyte, univariate exposure-response functions with the other five analytes held at their medians, and the overall mixture contrast between the joint 75th and joint 25th percentile profiles. PIPs describe selection support for a non-null contribution within the fitted surface; they are not effect sizes, they do not rank analytes by potency, and they are reported as exploratory throughout. Contrasts are on the scale of the modeled outcome, martingale residual or latent probit, and are not risk ratios. Because component-wise selection places a point mass at zero on the contrast whenever the posterior excludes every analyte, each model was also refitted with all six analytes forced into the surface, which gives an interval that does not mix a spike with a slab and so distinguishes a genuinely null surface from selection shrinkage. Four further summaries were computed from the forced-in fits. The overall exposure-response compares the surface with all six analytes at their qth percentile against all six at their 50th, for q from the 10th to the 90th. The single-analyte effect is the change in the surface when one analyte moves from its 25th to its 75th percentile while the other five are held at their 25th, 50th or 75th percentile. The single-analyte interaction summary is the difference between that effect with the others at their 75th percentile and with the others at their 25th, which departs from zero only if the response in one analyte depends on the level of the rest. The bivariate exposure-response gives the surface across the range of each analyte with a second analyte held at its 10th, 50th and 90th percentile and the remaining four at their medians, so that parallel curves indicate no interaction within that pair. Every summary is a difference between points on the same posterior draw, so the credible intervals account for the correlation between the points being compared.

Sensitivity analyses comprised omission of survey cycle; interview-based follow-up time; omission of smoking status; omission of the poverty–income ratio; serum cotinine in place of self-reported smoking; the extended model with alcohol use, physical activity and baseline disease; adjustment for serum albumin and the estimated glomerular filtration rate; restriction of the mixture to the four analytes with near-complete detection; quartile cut points assigned so that tied substituted values fall in the same quartile; adjustment for seafood consumption and total energy intake; and the design-based weighted models. Analyses were conducted in Python 3.13. Cox models were fitted with statsmodels and the weighted Cox models with lifelines; the Fine–Gray and kernel machine routines were implemented for this analysis. The kernel machine implementation was validated on simulated data with six correlated exposures, at the settings used here and pooling three chains per scenario. With two exposures given a true effect, one quadratic and one linear, posterior inclusion probabilities were 1.00 and 1.00 for the two active exposures and at most 0.0 for the four null exposures, and the 75th against 25th percentile contrast was recovered as 0.81 (0.72 to 0.92) against a true value of 0.85. With a single weak linear effect the active exposure was recovered with a posterior inclusion probability of 0.99 and a contrast of 0.20 (0.12 to 0.30) against a true value of 0.16, with at most 0.01 for the five null exposures. Under a null surface every posterior inclusion probability was 0.0 and the contrast was identically zero [41,42].

3. Results

3.1. Cohort, Follow-Up and Exposure

The primary sample comprised 5964 adults contributing 69,256 person-years, of whom 1119 (18.8 percent) died by 31 December 2019, including 344 cardiovascular and 235 cancer deaths. The crude mortality rate was 16.2 deaths per 1000 person-years. Median follow-up fell from 187 months in the 2003–2004 cycle to 119 months in 2009–2010, and the mortality rate fell from 19.1 to 13.4 per 1000 person-years while mean age was stable, so calendar period is associated with the rate of death as well as with the length of observation. Over the same period the legacy and emerging analytes moved in opposite directions: PFOS (perfluorooctane sulfonate) fell from a geometric mean of 20.04 to 9.33 ng/mL, while PFNA (perfluorononanoic acid) rose from 0.91 to 1.24 ng/mL (Table 1, Figure 1).

Table 1. Cohort, follow-up, mortality and exposure by survey cycle, in the primary analytic sample. GM, geometric mean in ng/mL; CVD, cardiovascular disease; PFOS, perfluorooctane sulfonate; PFNA, perfluorononanoic acid. Follow-up is person-months from examination to death or to the end of mortality linkage on 31 December 2019.

Cycle	*n*	Deaths	CVD Deaths	Cancer Deaths	Person-Years	Rate per 1000 py	Median Follow-Up, mo	Mean Age, y	PFOS GM	PFNA GM
2003–2004	1353	360	114	69	18,841	19.1	187	51.0	20.04	0.91
2005–2006	1395	276	89	59	17,705	15.6	165	48.2	16.68	1.06
2007–2008	1554	272	80	61	17,006	16.0	140	50.1	13.97	1.23
2009–2010	1662	211	61	46	15,704	13.4	119	49.8	9.33	1.24
Total	5964	1119	344	235	69,256	16.2	141	49.8	14.12	1.11

Serum concentrations spanned roughly two orders of magnitude within the cohort, with PFOS the most abundant analyte (geometric mean 14.12 ng/mL) and PFUnDA (perfluoroundecanoic acid) the least (0.20 ng/mL) (Table 2). The six analytes were moderately to strongly correlated on the logarithmic scale, with Spearman coefficients between 0.14 and 0.75 (median 0.50), the strongest pairing being PFNA with PFDeA. This degree of

correlation is what makes single-analyte estimates difficult to interpret and motivates the mixture approach.

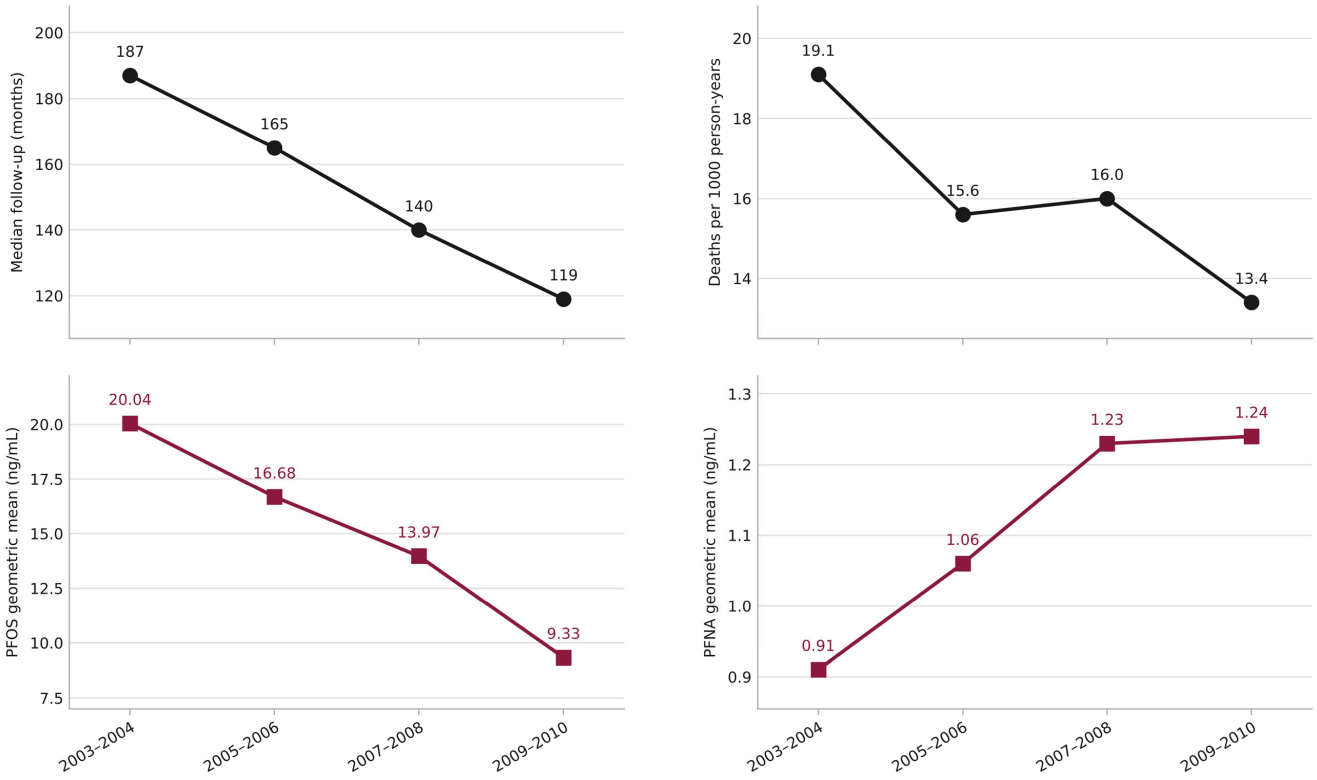

Figure 1. Median follow-up, all-cause mortality rate, and geometric mean PFOS (perfluorooctane sulfonate) and PFNA (perfluorononanoic acid) by NHANES (National Health and Nutrition Examination Survey) cycle, in the primary analytic sample. Follow-up and the mortality rate fall together across cycles, while PFOS falls and PFNA rises, so calendar period is associated with the exposure and with the length of observation at the same time. Note that each panel has its own vertical scale.

Table 2. Distribution of serum PFAS (per- and polyfluoroalkyl substance) concentrations in ng/mL in the primary analytic sample, the percentage above the limit of detection, and the rank correlation of each analyte with age. IQR, interquartile range. PFOS, perfluorooctane sulfonate; PFOA, perfluorooctanoic acid; PFNA, perfluorononanoic acid; PFHxS, perfluorohexane sulfonic acid; PFDeA, perfluorodecanoic acid; PFUnDA, perfluoroundecanoic acid.

Analyte	Geometric Mean	Median	IQR	5th–95th	Detected, %	Spearman rho with Age	*p*
PFOS	14.12	15.0	8.70–24.30	3.40–47.50	99.8	0.28	<0.001
PFOA	3.48	3.7	2.40–5.40	1.10–9.50	99.6	0.12	<0.001
PFNA	1.11	1.1	0.74–1.70	0.40–3.44	99.0	0.07	<0.001
PFHxS	1.67	1.8	1.00–3.00	0.30–7.20	98.2	0.17	<0.001
PFDeA	0.3	0.3	0.20–0.40	0.14–1.10	70.1	0.03	0.021
PFUnDA	0.2	0.2	0.14–0.20	0.07–0.80	60.2	0.09	<0.001

Every covariate used in any model was recorded for at least 89 percent of participants in every cycle, and for at least 92 percent of the analytic sample overall (Table 3). The family poverty–income ratio, at 92.8 percent, was the least complete and determines the size of the primary sample.

Table 3. Availability of covariates by survey cycle, among the 6431 participants with complete serum PFAS (per- and polyfluoroalkyl substances) who met the age and linkage criteria. Percentages are of participants with a non-missing value. PSU, primary sampling unit.

Variable	2003–2004	2005–2006	2007–2008	2009–2010	Analytic Sample (n = 6431)
Age	100.0%	100.0%	100.0%	100.0%	100.0%
Sex	100.0%	100.0%	100.0%	100.0%	100.0%
Race/Hispanic origin	100.0%	100.0%	100.0%	100.0%	100.0%
Education	100.0%	100.0%	100.0%	100.0%	100.0%
Body mass index	100.0%	100.0%	100.0%	100.0%	100.0%
Smoking status	100.0%	100.0%	99.9%	100.0%	100.0%
Serum cotinine	100.0%	100.0%	100.0%	99.9%	100.0%
Household income	94.7%	99.0%	99.6%	99.4%	98.3%
Poverty–income ratio	95.1%	96.1%	90.7%	90.2%	92.8%
Alcohol use	92.5%	93.4%	94.2%	89.3%	92.2%
Leisure-time physical activity	99.9%	100.0%	100.0%	100.0%	100.0%
Self-reported cardiovascular disease	99.6%	99.7%	99.4%	99.5%	99.5%
Self-reported cancer	99.9%	99.9%	99.9%	100.0%	99.9%
Self-reported diabetes	100.0%	99.9%	99.8%	99.8%	99.9%
Self-reported hypertension	98.9%	99.7%	99.9%	99.8%	99.6%
Serum albumin	99.6%	99.8%	99.9%	100.0%	99.9%
Serum creatinine	99.6%	99.8%	99.9%	100.0%	99.8%
PFAS subsample weight	100.0%	100.0%	100.0%	100.0%	100.0%
Survey stratum	100.0%	100.0%	100.0%	100.0%	100.0%
Survey PSU	100.0%	100.0%	100.0%	100.0%	100.0%

3.2. Joint Mixture Association

A simultaneous one-quartile increase in all six PFAS (per- and polyfluoroalkyl substances) was associated with lower all-cause mortality, hazard ratio (HR) 0.83 (95 percent confidence interval (CI) 0.77 to 0.90, $p < 0.001$), and with 0.81 (0.75 to 0.88) per doubling of every analyte. There was no clear evidence of an association with cardiovascular mortality, 0.92 (0.79, 1.07), or with cancer mortality, 0.90 (0.75, 1.08) (Table 4). Fine–Gray subdistribution models, which account for competing causes of death, gave 0.99 (0.85, 1.17) for cardiovascular and 0.97 (0.80, 1.19) for cancer mortality; the cause-specific and subdistribution estimates differ, as expected when a substantial fraction of the cohort dies of other causes, and neither excludes the null. The design-based estimates, using the PFAS (per- and polyfluoroalkyl substances) subsample weights with the variance clustered on the primary

sampling unit, were 0.82 (0.74, 0.90) for all-cause, 0.90 (0.75, 1.08) for cardiovascular and 0.87 (0.71, 1.08) for cancer mortality, so the design-based and unweighted analyses agree on every endpoint (Table 4).

Table 4. Cox proportional hazards estimates for each analyte and for the joint mixture, adjusted for age, sex, race and Hispanic origin, education, body mass index, survey cycle, smoking status and the family poverty–income ratio. Cause-specific models censor deaths from other causes. The survey-weighted rows use the PFAS (per- and polyfluoroalkyl substances) subsample weights with the variance clustered on the primary sampling unit within survey stratum. CVD, cardiovascular disease; CI, confidence interval.

Outcome	Scale	Exposure	*n*	Deaths	Hazard Ratio (95% CI)	*p*
All-cause	per quartile	PFOS	5964	1119	0.98 (0.90, 1.07)	0.679
All-cause	per quartile	PFOA	5964	1119	0.94 (0.87, 1.01)	0.108
All-cause	per quartile	PFNA	5964	1119	0.95 (0.87, 1.03)	0.226
All-cause	per quartile	PFHxS	5964	1119	1.00 (0.93, 1.07)	0.922
All-cause	per quartile	PFDeA	5964	1119	0.96 (0.89, 1.05)	0.369
All-cause	per quartile	PFUnDA	5964	1119	0.99 (0.92, 1.06)	0.791
All-cause	per quartile	Joint (all six)	5964	1119	0.83 (0.77, 0.90)	<0.001
All-cause	per quartile	Joint (all six), survey-weighted	5964	1119	0.82 (0.74, 0.90)	<0.001
All-cause	per doubling	PFOS	5964	1119	0.98 (0.89, 1.07)	0.620
All-cause	per doubling	PFOA	5964	1119	0.87 (0.79, 0.95)	0.003
All-cause	per doubling	PFNA	5964	1119	0.96 (0.88, 1.05)	0.378
All-cause	per doubling	PFHxS	5964	1119	1.00 (0.94, 1.07)	0.955
All-cause	per doubling	PFDeA	5964	1119	1.00 (0.90, 1.11)	0.977
All-cause	per doubling	PFUnDA	5964	1119	1.00 (0.91, 1.10)	0.956
All-cause	per doubling	Joint (all six)	5964	1119	0.81 (0.75, 0.88)	<0.001
All-cause	per doubling	Joint (all six), survey-weighted	5964	1119	0.83 (0.76, 0.92)	<0.001
CVD	per quartile	PFOS	5964	344	1.10 (0.94, 1.28)	0.253
CVD	per quartile	PFOA	5964	344	0.87 (0.76, 0.99)	0.038
CVD	per quartile	PFNA	5964	344	0.99 (0.85, 1.16)	0.895
CVD	per quartile	PFHxS	5964	344	0.96 (0.85, 1.08)	0.524
CVD	per quartile	PFDeA	5964	344	0.96 (0.83, 1.12)	0.638
CVD	per quartile	PFUnDA	5964	344	1.05 (0.93, 1.19)	0.453
CVD	per quartile	Joint (all six)	5964	344	0.92 (0.79, 1.07)	0.262
CVD	per quartile	Joint (all six), survey-weighted	5964	344	0.90 (0.75, 1.08)	0.249
CVD	per doubling	PFOS	5964	344	1.01 (0.85, 1.20)	0.913
CVD	per doubling	PFOA	5964	344	0.82 (0.69, 0.96)	0.015
CVD	per doubling	PFNA	5964	344	1.11 (0.94, 1.31)	0.217
CVD	per doubling	PFHxS	5964	344	0.95 (0.85, 1.07)	0.399
CVD	per doubling	PFDeA	5964	344	0.96 (0.80, 1.16)	0.680

Table 4. *Cont.*

Outcome	Scale	Exposure	*n*	Deaths	Hazard Ratio (95% CI)	*p*
CVD	per doubling	PFUnDA	5964	344	1.05 (0.89, 1.24)	0.561
CVD	per doubling	Joint (all six)	5964	344	0.88 (0.76, 1.01)	0.075
CVD	per doubling	Joint (all six), survey-weighted	5964	344	0.89 (0.75, 1.06)	0.190
Cancer	per quartile	PFOS	5964	235	0.88 (0.73, 1.07)	0.202
Cancer	per quartile	PFOA	5964	235	1.05 (0.89, 1.23)	0.587
Cancer	per quartile	PFNA	5964	235	0.95 (0.79, 1.15)	0.600
Cancer	per quartile	PFHxS	5964	235	1.04 (0.90, 1.21)	0.562
Cancer	per quartile	PFDeA	5964	235	1.03 (0.86, 1.24)	0.734
Cancer	per quartile	PFUnDA	5964	235	0.95 (0.82, 1.11)	0.530
Cancer	per quartile	Joint (all six)	5964	235	0.90 (0.75, 1.08)	0.269
Cancer	per quartile	Joint (all six), survey-weighted	5964	235	0.87 (0.71, 1.08)	0.203
Cancer	per doubling	PFOS	5964	235	0.96 (0.78, 1.17)	0.663
Cancer	per doubling	PFOA	5964	235	1.02 (0.84, 1.26)	0.814
Cancer	per doubling	PFNA	5964	235	0.87 (0.72, 1.06)	0.181
Cancer	per doubling	PFHxS	5964	235	1.01 (0.88, 1.16)	0.879
Cancer	per doubling	PFDeA	5964	235	1.08 (0.86, 1.36)	0.500
Cancer	per doubling	PFUnDA	5964	235	1.03 (0.84, 1.25)	0.796
Cancer	per doubling	Joint (all six)	5964	235	0.96 (0.81, 1.14)	0.671
Cancer	per doubling	Joint (all six), survey-weighted	5964	235	0.97 (0.79, 1.19)	0.760

In the quantile g-computation decomposition for all-cause mortality every analyte coefficient was negative, so the negative partial sum was −0.183 against a positive partial sum of 0.000 and the mixture acts entirely in one direction. The shares of that negative sum were PFOA (weight 0.33), PFNA (weight 0.29), PFDeA (weight 0.21), PFOS (weight 0.10), PFUnDA (weight 0.05), PFHxS (weight 0.02). For cardiovascular and cancer mortality the positive and negative partial sums were of similar magnitude and largely canceled, which is the decomposition of a null joint estimate rather than evidence of offsetting analyte effects. Among individual analytes only PFOA (perfluorooctanoic acid) was associated with all-cause mortality, 0.87 (0.79, 0.95) per doubling, and with cardiovascular mortality, 0.82 (0.69, 0.96) per doubling and 0.87 (0.76, 0.99) per quartile; no analyte was associated with cancer mortality.

The proportional hazards assumption was tested for each exposure in each of the three models, eighteen tests in total. Two of the tests gave *p*-values below 0.05: PFOS for cvd mortality ($p = 0.045$); PFDeA for cancer mortality ($p = 0.021$). This is close to the number expected by chance at that threshold, and the period-specific analysis in Section 3.3 addresses the substantive form of non-proportionality directly.

3.3. Lag and Period-Specific Analyses

In the landmark analysis at two years, restricted to the 5836 participants alive and under observation at 24 months with time measured from the landmark, the all-cause estimate was 0.84 (0.77, 0.92). At the five-year landmark, restricted to 5582 participants, it

attenuated to 0.88 (0.79, 0.97), a shift of about a quarter of the association on the logarithmic scale. The simpler construction that deletes early deaths while still measuring follow-up from examination gave the same estimates to two decimal places, so the immortal person-time it carries does not materially change the estimate in these data (Table 5, Figure 2). The period-specific analysis addresses the same question from a different direction: the joint hazard ratio (HR) was 0.80 (0.73, 0.88) for deaths occurring in the first ten years of follow-up and 0.96 (0.80, 1.16) for deaths occurring thereafter. Because the two periods use disjoint risk sets the partial likelihood factorizes across them, so the two estimates are independent and their difference can be tested directly: the ratio of the early to the late hazard ratio was 0.83 (0.67 to 1.02, $p = 0.070$). The inverse association is therefore larger in early follow-up and close to null among longer-term survivors, although the difference between periods is itself imprecisely estimated and does not reach conventional significance. The corresponding ratios were 0.72 (0.49 to 1.07, $p = 0.108$) for cardiovascular and 0.94 (0.61 to 1.47, $p = 0.797$) for cancer mortality.

Table 5. Joint mixture hazard ratios for a one-quartile increase in all six PFAS (per- and polyfluoroalkyl substances) under lag analyses and within periods of follow-up. Landmark models restrict the risk set to participants alive and under observation at the landmark and measure time from it, so that person-time and deaths before the landmark do not contribute. The rows that remove early deaths while measuring follow-up from examination are shown for comparison and carry immortal person-time. Period-specific models censor participants at ten years or admit them at ten years respectively. CVD, cardiovascular disease; CI, confidence interval.

Outcome	Analysis	*n*	Deaths	Joint Hazard Ratio (95% CI)	*p*
All-cause	Primary, all deaths	5964	1119	0.83 (0.77, 0.90)	<0.001
All-cause	Landmark at 2 years, time from landmark	5836	991	0.84 (0.77, 0.92)	<0.001
All-cause	Landmark at 5 years, time from landmark	5582	737	0.88 (0.79, 0.97)	0.013
All-cause	Early deaths within 2 years removed, follow-up from examination	5836	991	0.84 (0.77, 0.92)	<0.001
All-cause	Early deaths within 5 years removed, follow-up from examination	5582	737	0.88 (0.79, 0.97)	0.013
All-cause	Deaths in the first 10 years of follow-up only	5964	850	0.80 (0.73, 0.88)	<0.001
All-cause	Deaths after 10 years of follow-up only	4376	269	0.96 (0.80, 1.16)	0.701
All-cause	Ratio of early to late hazard ratio	5964	1119	0.83 (0.67, 1.02)	0.070
CVD	Primary, all deaths	5964	344	0.92 (0.79, 1.07)	0.262
CVD	Landmark at 2 years, time from landmark	5836	302	0.94 (0.80, 1.10)	0.414
CVD	Landmark at 5 years, time from landmark	5582	217	0.96 (0.79, 1.17)	0.711
CVD	Early deaths within 2 years removed, follow-up from examination	5922	302	0.94 (0.80, 1.10)	0.414
CVD	Early deaths within 5 years removed, follow-up from examination	5837	217	0.96 (0.79, 1.17)	0.711
CVD	Deaths in the first 10 years of follow-up only	5964	263	0.87 (0.74, 1.03)	0.099
CVD	Deaths after 10 years of follow-up only	4376	81	1.20 (0.84, 1.71)	0.316
CVD	Ratio of early to late hazard ratio	5964	344	0.72 (0.49, 1.07)	0.108
Cancer	Primary, all deaths	5964	235	0.90 (0.75, 1.08)	0.269
Cancer	Landmark at 2 years, time from landmark	5836	205	0.90 (0.74, 1.09)	0.286

Table 5. *Cont.*

Outcome	Analysis	*n*	Deaths	Joint Hazard Ratio (95% CI)	*p*
Cancer	Landmark at 5 years, time from landmark	5582	153	0.93 (0.74, 1.17)	0.541
Cancer	Early deaths within 2 years removed, follow-up from examination	5934	205	0.90 (0.74, 1.09)	0.286
Cancer	Early deaths within 5 years removed, follow-up from examination	5882	153	0.93 (0.74, 1.17)	0.541
Cancer	Deaths in the first 10 years of follow-up only	5964	179	0.88 (0.72, 1.08)	0.237
Cancer	Deaths after 10 years of follow-up only	4376	56	0.94 (0.63, 1.39)	0.747
Cancer	Ratio of early to late hazard ratio	5964	235	0.94 (0.61, 1.47)	0.797

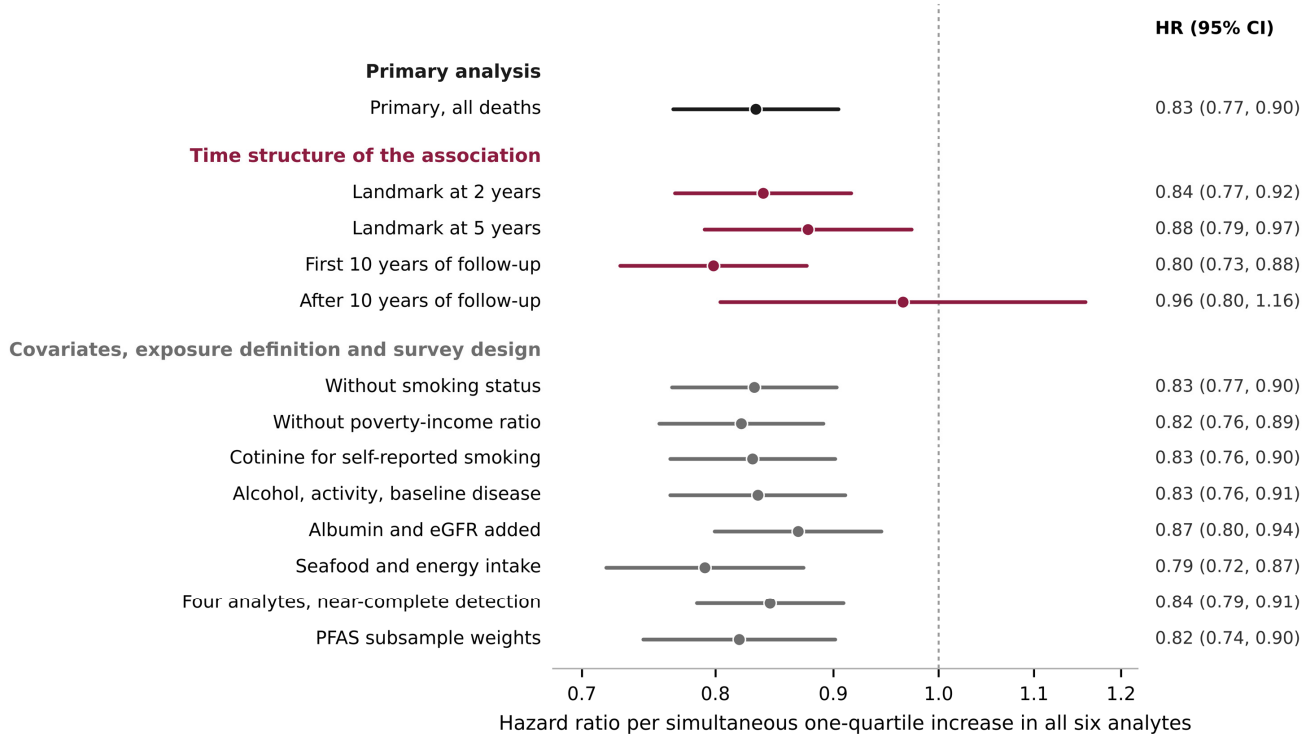

Figure 2. Joint all-cause mixture hazard ratio (HR) for a one-quartile increase in all six PFAS (per- and polyfluoroalkyl substances), with 95 percent confidence intervals (CI). Estimates are grouped into three families, and the color of each estimate matches the color of its heading: the primary analysis in black, the analyses of the time structure of the association in red, and the covariate, exposure-definition and survey-design analyses in gray. The dashed line marks the null. The horizontal axis is on a logarithmic scale. eGFR, estimated glomerular filtration rate.

3.4. Sensitivity Analyses

The all-cause estimate was stable across model specification (Table 6). Omitting survey cycle, using interview-based follow-up time, omitting smoking status, omitting the poverty–income ratio (PIR), substituting serum cotinine for self-reported smoking, and adding alcohol use, physical activity and baseline disease all left it between 0.82 and 0.83. Restricting the mixture to the four analytes with near-complete detection gave 0.84 (0.79, 0.91), and assigning quartile cut points so that tied substituted values fall in the same quartile gave 0.83 (0.77, 0.90). The design-based model using the PFAS (per- and polyfluoroalkyl substances) subsample weights, with the variance clustered on the primary

sampling unit, gave 0.82 (0.74, 0.90), so neither the weights nor the clustering changes the conclusion. The single specification that moved the estimate was adjustment for serum albumin and the estimated glomerular filtration rate (eGFR), which attenuated it to 0.87 (0.80, 0.94), about a quarter of the association on the logarithmic scale.

Table 6. Sensitivity analyses for the joint mixture hazard ratio per one-quartile increase in all six analytes, by outcome. CVD, cardiovascular disease; CI, confidence interval; PIR, poverty–income ratio; eGFR, estimated glomerular filtration rate; PFAS, per- and polyfluoroalkyl substances.

Outcome	Analysis	*n*	Deaths	Joint Hazard Ratio (95% CI)	*p*
All-cause	Primary model	5964	1119	0.83 (0.77, 0.90)	<0.001
All-cause	Without adjustment for survey cycle	5964	1119	0.84 (0.77, 0.91)	<0.001
All-cause	Interview-based follow-up time	5964	1119	0.83 (0.77, 0.90)	<0.001
All-cause	Original covariate set, full sample (no smoking or PIR)	6431	1196	0.83 (0.77, 0.90)	<0.001
All-cause	Without smoking status	5964	1119	0.83 (0.77, 0.90)	<0.001
All-cause	Without poverty–income ratio	5964	1119	0.82 (0.76, 0.89)	<0.001
All-cause	Serum cotinine in place of self-reported smoking	5963	1118	0.83 (0.76, 0.90)	<0.001
All-cause	Extended: alcohol, physical activity and baseline disease	5477	1021	0.83 (0.76, 0.91)	<0.001
All-cause	Serum albumin and estimated glomerular filtration rate added	5954	1117	0.87 (0.80, 0.94)	<0.001
All-cause	Four analytes with near-complete detection	5964	1119	0.84 (0.79, 0.91)	<0.001
All-cause	PFAS subsample weights, variance clustered on primary sampling unit	5964	1119	0.82 (0.74, 0.90)	<0.001
All-cause	Participants with dietary data, without diet	4765	729	0.80 (0.72, 0.88)	<0.001
All-cause	Adjusted for seafood consumption	4765	729	0.79 (0.72, 0.87)	<0.001
All-cause	Adjusted for seafood consumption and energy intake	4765	729	0.79 (0.72, 0.87)	<0.001
CVD	Primary model	5964	344	0.92 (0.79, 1.07)	0.262
CVD	Without adjustment for survey cycle	5964	344	0.92 (0.79, 1.07)	0.268
CVD	Interview-based follow-up time	5964	344	0.92 (0.79, 1.07)	0.262
CVD	Original covariate set, full sample (no smoking or PIR)	6431	372	0.92 (0.80, 1.06)	0.261
CVD	Without smoking status	5964	344	0.92 (0.79, 1.07)	0.261
CVD	Without poverty–income ratio	5964	344	0.90 (0.78, 1.05)	0.172
CVD	Serum cotinine in place of self-reported smoking	5963	344	0.92 (0.79, 1.06)	0.255
CVD	Extended: alcohol, physical activity and baseline disease	5477	321	0.89 (0.76, 1.04)	0.158
CVD	Serum albumin and estimated glomerular filtration rate added	5954	342	0.94 (0.80, 1.09)	0.390

Table 6. *Cont.*

Outcome	Analysis	*n*	Deaths	Joint Hazard Ratio (95% CI)	*p*
CVD	Four analytes with near-complete detection	5964	344	0.90 (0.79, 1.03)	0.122
CVD	PFAS subsample weights, variance clustered on primary sampling unit	5964	344	0.90 (0.75, 1.08)	0.249
CVD	Participants with dietary data, without diet	4765	217	0.87 (0.73, 1.04)	0.126
CVD	Adjusted for seafood consumption	4765	217	0.85 (0.71, 1.02)	0.081
CVD	Adjusted for seafood consumption and energy intake	4765	217	0.85 (0.71, 1.02)	0.079
CVD	Fine–Gray subdistribution hazard (competing risks)	5964	344	0.99 (0.85, 1.17)	0.927
Cancer	Primary model	5964	235	0.90 (0.75, 1.08)	0.269
Cancer	Without adjustment for survey cycle	5964	235	0.91 (0.76, 1.09)	0.301
Cancer	Interview-based follow-up time	5964	235	0.90 (0.76, 1.08)	0.272
Cancer	Original covariate set, full sample (no smoking or PIR)	6431	252	0.88 (0.74, 1.04)	0.140
Cancer	Without smoking status	5964	235	0.90 (0.75, 1.07)	0.234
Cancer	Without poverty–income ratio	5964	235	0.90 (0.75, 1.08)	0.249
Cancer	Serum cotinine in place of self-reported smoking	5963	235	0.89 (0.75, 1.07)	0.210
Cancer	Extended: alcohol, physical activity and baseline disease	5477	219	0.95 (0.78, 1.14)	0.562
Cancer	Serum albumin and estimated glomerular filtration rate added	5954	235	0.93 (0.78, 1.12)	0.461
Cancer	Four analytes with near-complete detection	5964	235	0.92 (0.78, 1.08)	0.319
Cancer	PFAS subsample weights, variance clustered on primary sampling unit	5964	235	0.87 (0.71, 1.08)	0.203
Cancer	Participants with dietary data, without diet	4765	165	0.90 (0.73, 1.10)	0.311
Cancer	Adjusted for seafood consumption	4765	165	0.90 (0.73, 1.11)	0.309
Cancer	Adjusted for seafood consumption and energy intake	4765	165	0.91 (0.73, 1.12)	0.358
Cancer	Fine–Gray subdistribution hazard (competing risks)	5964	235	0.97 (0.80, 1.19)	0.760

Socioeconomic position behaved in two ways that need to be separated. The crude configuration is exactly the one that produces a spurious inverse association: across household income brackets the mean mixture quartile score rose monotonically from 8.43 in those earning under $20,000 to 9.92 in those earning $75,000 or more, while the proportion who died fell from 29.9 to 8.0 percent. Adjustment, however, did not attenuate the mixture estimate. Among the 5928 participants with household income recorded, the all-cause estimate was 0.82 (0.75 to 0.89) without income and 0.83 (0.76 to 0.90) with it, and omitting the PIR (poverty–income ratio) from the primary model moved the estimate from 0.83 to

0.82, that is, slightly away from the null rather than towards it. Measured socioeconomic position therefore has the structure that could generate the association but does not account for it.

Dietary intake behaved in the same way. Across categories of seafood consumption in the past 30 days the mean mixture quartile score rose from 7.32 among those reporting no fish or shellfish meals to 9.92 among those reporting six or more, mean serum PFOS (perfluorooctane sulfonate) from 14.6 to 18.9 ng/mL, and the mean poverty–income ratio (PIR) from 2.05 to 2.95, while the proportion who died was lower in the highest category than the lowest (14.2 against 17.0 percent). Seafood consumption therefore tracks both higher exposure and higher socioeconomic position, as expected. Adjusting for it did not attenuate the mixture estimate either: among the 4765 participants with dietary data the all-cause estimate was 0.80 (0.72, 0.88) without the dietary variables and 0.79 (0.72, 0.87) with seafood consumption and total energy intake.

3.5. Exposure-Response Surface

The exploratory Bayesian kernel machine regression (BKMR) analysis, fitted to martingale residuals from a covariate-only Cox model, was consistent with the regression findings. The overall mixture contrast between the joint 75th and joint 25th percentile profiles was −0.023 (−0.040, −0.007) for all-cause mortality, with posterior probability 0.003 of exceeding zero, and PFOA (perfluorooctanoic acid) was the only analyte with appreciable posterior selection support (posterior inclusion probability (PIP) 0.902, against 0.105 for PFNA (perfluorononanoic acid) and 0.01 or below for the remaining four); its univariate exposure-response function declined across the exposure range. The contrast was +0.000 (−0.006, +0.007) for cardiovascular and −0.000 (+0.000, +0.000) for cancer mortality (Table 7, Figure 3).

Table 7. Bayesian kernel machine regression posterior inclusion probabilities and the overall mixture contrast between the joint 75th and joint 25th exposure percentiles, pooled across chains. Contrasts are on the scale of the modeled outcome and are not risk ratios. CVD, cardiovascular disease; PFOS, perfluorooctane sulfonate; PFOA, perfluorooctanoic acid; PFNA, perfluorononanoic acid; PFHxS, perfluorohexane sulfonic acid; PFDeA, perfluorodecanoic acid; PFUnDA, perfluoroundecanoic acid.

Family	Outcome	Chains	PFOS	PFOA	PFNA	PFHxS	PFDeA	PFUnDA	Overall Contrast (75th vs. 25th)	Posterior p (>0)
Probit scale, selection	All-cause	3	0.001	1.0	0.006	0.001	0.01	0.002	−0.114 (−0.190, −0.037)	0.003
Probit scale, selection	CVD	3	0.044	0.006	0.002	0.002	0.024	0.001	−0.002 (−0.055, +0.000)	0.024
Probit scale, selection	Cancer	3	0.018	0.037	0.037	0.023	0.004	0.035	−0.003 (−0.016, +0.000)	0.02
Residual scale, forced in	All-cause	3	forced in	forced in	forced in	forced in	forced in	forced in	−0.033 (−0.055, −0.012)	0.001
Residual scale, forced in	CVD	3	forced in	forced in	forced in	forced in	forced in	forced in	−0.003 (−0.015, +0.009)	0.336
Residual scale, forced in	Cancer	3	forced in	forced in	forced in	forced in	forced in	forced in	−0.005 (−0.015, +0.004)	0.152
Residual scale, selection	All-cause	3	0.001	0.902	0.105	0.0	0.001	0.001	−0.023 (−0.040, −0.007)	0.003

Table 7. *Cont.*

Family	Outcome	Chains	PFOS	PFOA	PFNA	PFHxS	PFDeA	PFUnDA	Overall Contrast (75th vs. 25th)	Posterior p (>0)
Residual scale, selection	CVD	3	0.163	0.031	0.002	0.002	0.01	0.001	+0.000 (−0.006, +0.007)	0.1
Residual scale, selection	Cancer	3	0.0	0.001	0.013	0.0	0.0	0.002	−0.000 (+0.000, +0.000)	0.002

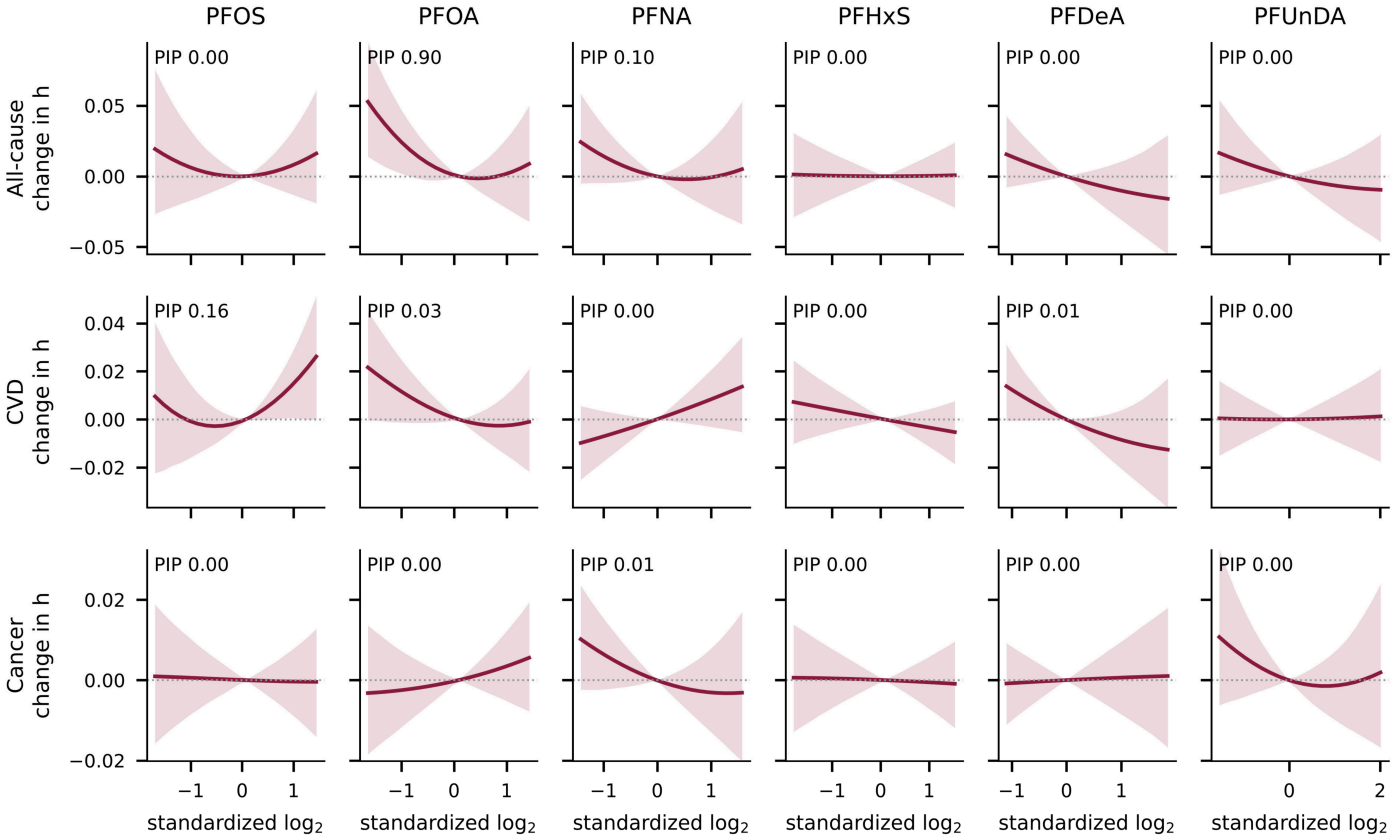

Figure 3. Bayesian kernel machine regression univariate exposure-response functions for each analyte with the other five held at their medians, fitted to martingale residuals from a covariate-only Cox model with every analyte forced into the surface, so that the bands show the uncertainty in the exposure-response function rather than the posterior probability of exclusion. In each panel the solid red curve is the posterior mean exposure-response function, the shaded band around it is the 95 percent credible interval, and the dotted horizontal line marks no change in the log hazard. Posterior inclusion probabilities from the corresponding variable-selection models are shown in each panel. Vertical scales are common within each outcome row. CVD, cardiovascular disease.

The cause-specific contrasts are very close to zero with narrow intervals, and this requires explanation rather than acceptance at face value. It is a property of component-wise variable selection, not of the numerical implementation. When the data provide little support for any analyte, the posterior concentrates on the model in which every indicator is zero; in that model the exposure surface is flat and the contrast is exactly zero, so the posterior for the contrast is a point mass at zero mixed with a small slab from the minority of iterations in which some analyte is included. The resulting interval is narrow because exclusion is near-certain, not because a small effect has been estimated precisely. To separate the two, the models were refitted with every analyte forced into the surface. The contrast

was then −0.003 (−0.015, +0.009) for cardiovascular and −0.005 (−0.015, +0.004) for cancer mortality, against −0.033 (−0.055, −0.012) for all-cause mortality (Table 7). The cause-specific intervals widen and cover zero, which is the honest summary: these data do not identify a mixture effect on cardiovascular or cancer mortality, and they do not rule out a modest one either.

The probit specification on the binary outcome, which corresponds to the analysis reported in the earlier mixture literature, gave −0.114 (−0.190, −0.037) for all-cause mortality and −0.002 (−0.055, +0.000) and −0.003 (−0.016, +0.000) for cardiovascular and cancer mortality. It agrees with the residual-scale analysis in direction, in which analyte carries selection support and in which endpoints show no association, but it required more than twice as many iterations to converge, because replacing the follow-up time with a binary indicator removes the information that distinguishes these cycles. That is the point made in Section 4, and it is the reason the residual-scale models rather than the probit models are reported as the exploratory analysis of the surface.

Mixing was satisfactory for the quantities that carry the reported results. Across all nine models the split Gelman–Rubin statistic for the overall mixture contrast was at most 1.016 and its effective sample size at least 140; for the primary all-cause model the values were 1.002 and 1388. Chain-specific posterior inclusion probabilities for PFOA in that model were 0.87, 0.99, 0.84 and chain-specific contrasts −0.024, −0.023, −0.024, so the posterior is not sensitive to the starting seed. Across every reported parameter of every model, the probit models included, the largest defined statistic was 1.074. Two features of Table S1 of the Supplementary Materials need stating. The statistic is not defined for the PFHxS (perfluorohexane sulfonic acid) kernel coefficient in the all-cause selection model, because that analyte was excluded in every retained draw of all three chains, so its coefficient is identically zero and the within-chain variance the statistic divides by is zero; this is unanimous agreement between chains rather than a failure to converge. Separately, 9 of the 72 reported parameters have effective sample sizes below 100, between 46 and 91. All of them are analyte-specific kernel coefficients rather than overall mixture contrasts, and 8 of the 9 are in the cancer models, where 235 deaths support the surface; the remaining one is a kernel coefficient in the cardiovascular forced-in model. Those coefficients are nuisance parameters that govern the smoothness of the surface rather than quantities reported here, but their limited mixing is a further reason to treat the cancer surface as exploratory.

3.6. Overall, Single-Analyte and Interaction Summaries

The four standard kernel machine summaries were computed from the forced-in fits, for the reason given above: with component-wise selection the cause-specific surfaces collapse onto a point mass at zero and every summary is identically zero, which describes the selection prior rather than the exposure.

The overall exposure-response, which moves all six analytes together away from their joint median, declined monotonically across the whole exposure range for all-cause mortality. Relative to all six at their joint 50th percentile, the surface was +0.056 (+0.026, +0.086) with all six at their joint 10th percentile and −0.022 (−0.049, +0.004) at their joint 90th. The cardiovascular and cancer curves showed no consistent trend, running from +0.007 (−0.009, +0.024) to −0.004 (−0.020, +0.011) and from +0.015 (+0.002, +0.028) to +0.005 (−0.007, +0.017), respectively, between the joint 10th and 90th percentiles (Figure 4).

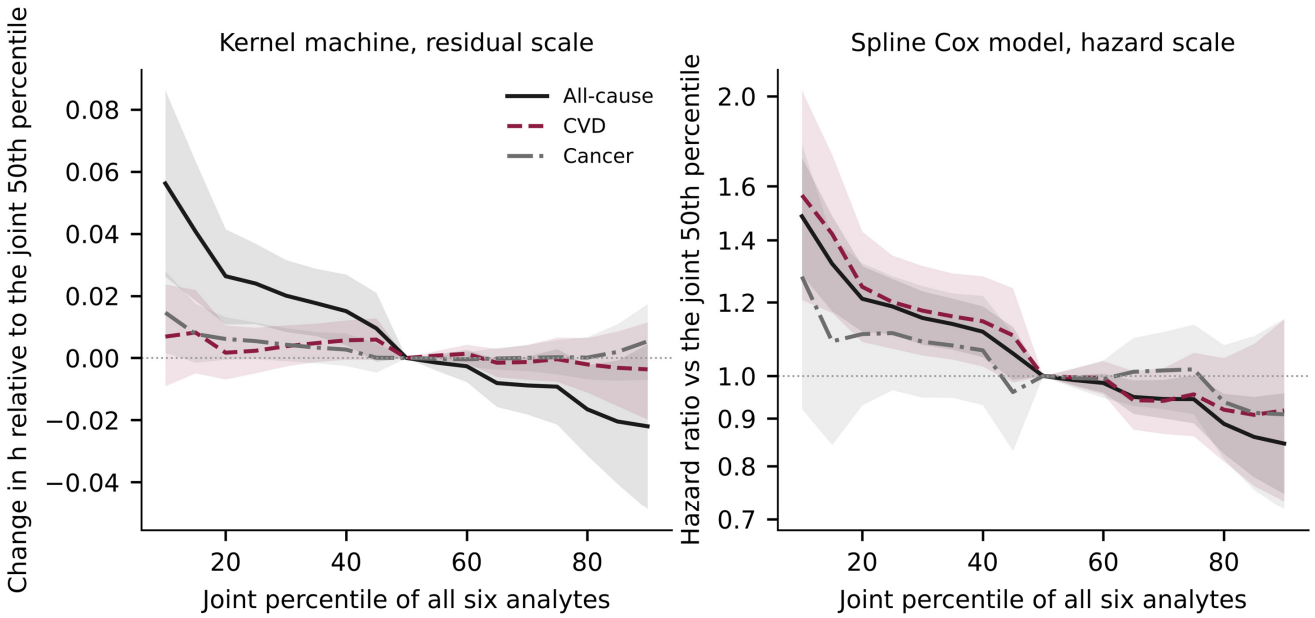

Figure 4. Joint exposure-response when all six PFAS (per- and polyfluoroalkyl substances) are moved together to the stated joint percentile, relative to all six at their joint 50th percentile, with 95 percent intervals. (**Left**) the Bayesian kernel machine regression surface on the residual scale. (**Right**) the hazard ratio from the spline Cox model. All-cause mortality is drawn with a solid black line, cardiovascular mortality with a dashed red line and cancer mortality with a dash-dotted gray line, and the shaded bands are the 95 percent intervals. The dotted horizontal line marks the joint 50th percentile reference, at which the change in h is zero and the hazard ratio is one. CVD, cardiovascular disease.

Single-analyte effects, the change in the surface when one analyte moves from its 25th to its 75th percentile with the other five held fixed, did not exclude zero for any analyte at any endpoint with the other five at their medians (Table 8, Figure S1). This is the expected consequence of forcing six correlated analytes into one surface: the joint movement of the mixture is identified while the contribution of any single analyte with the others held fixed is not. The variable-selection fits resolve the same information differently, by concentrating posterior support on PFOA (perfluorooctanoic acid), and neither result licenses a claim about which analyte carries the association.

Interaction summaries, the difference between each single-analyte effect with the other five at their 75th percentile and with the other five at their 25th, covered zero for 16 of the 18 analyte and endpoint combinations. The exceptions were PFOS for cvd mortality, −0.027 (−0.055, −0.002); and PFOA for cvd mortality, −0.023 (−0.043, −0.003). With 18 intervals examined and no adjustment for multiplicity, and with the component estimates that make up each difference themselves covering zero in most cases, these are hypotheses rather than findings. The bivariate exposure-response functions locate the pattern in a single pair. For all-cause mortality the largest separation between the curves at the 10th and 90th percentiles of the second analyte was 0.08, for PFOS with PFOA, against 0.03 or less for every other ordered pair (Figure 5). The two analytes concerned are the two most abundant in the panel and the second most strongly correlated pair of the fifteen, with a Spearman coefficient of 0.64. Higher PFOS was associated with a higher surface only among participants with low PFOA. Given the number of comparisons, the width of the intervals and the sparseness of that corner of the exposure space, this is a pattern to test in an independent sample rather than a result to interpret here. For the remaining pairs there is no evidence that the association of any analyte with mortality depends on the level of

the others, and the joint estimates reported above can be read as additive on the scale of the surface.

Table 8. Bayesian kernel machine regression single-analyte effects and interaction summaries, from the fits with every analyte forced into the surface. The effect is the change in the surface when the analyte moves from its 25th to its 75th percentile with the other five held at the stated percentile. The interaction is the effect with the others at their 75th percentile minus the effect with the others at their 25th. The single-analyte effects are drawn as Figure S1 of the Supplementary Materials. Values are on the scale of the modeled outcome and are not risk ratios. CVD, cardiovascular disease.

Outcome	Analyte	Effect, Others at 25th	Effect, Others at 50th	Effect, Others at 75th	Interaction
All-cause	PFOS	+0.014 (−0.017, +0.051)	+0.002 (−0.021, +0.025)	−0.010 (−0.042, +0.019)	−0.023 (−0.074, +0.009)
All-cause	PFOA	+0.001 (−0.032, +0.036)	−0.011 (−0.034, +0.013)	−0.021 (−0.050, +0.005)	−0.022 (−0.063, +0.010)
All-cause	PFNA	−0.010 (−0.037, +0.018)	−0.009 (−0.032, +0.015)	−0.007 (−0.032, +0.019)	+0.003 (−0.021, +0.027)
All-cause	PFHxS	+0.001 (−0.016, +0.020)	−0.000 (−0.017, +0.017)	−0.001 (−0.021, +0.018)	−0.003 (−0.022, +0.012)
All-cause	PFDeA	−0.012 (−0.036, +0.009)	−0.012 (−0.033, +0.008)	−0.012 (−0.035, +0.009)	+0.000 (−0.019, +0.017)
All-cause	PFUnDA	−0.005 (−0.015, +0.006)	−0.005 (−0.015, +0.004)	−0.005 (−0.016, +0.005)	−0.000 (−0.010, +0.009)
CVD	PFOS	+0.025 (+0.004, +0.045)	+0.011 (−0.004, +0.024)	−0.003 (−0.021, +0.015)	−0.027 (−0.055, −0.002)
CVD	PFOA	+0.004 (−0.013, +0.021)	−0.008 (−0.021, +0.005)	−0.019 (−0.035, −0.004)	−0.023 (−0.043, −0.003)
CVD	PFNA	+0.007 (−0.007, +0.020)	+0.009 (−0.003, +0.021)	+0.011 (−0.002, +0.025)	+0.004 (−0.006, +0.018)
CVD	PFHxS	−0.005 (−0.016, +0.005)	−0.004 (−0.014, +0.005)	−0.003 (−0.014, +0.008)	+0.002 (−0.008, +0.014)
CVD	PFDeA	−0.013 (−0.027, +0.001)	−0.011 (−0.022, +0.001)	−0.008 (−0.022, +0.004)	+0.004 (−0.010, +0.017)
CVD	PFUnDA	+0.000 (−0.005, +0.006)	−0.000 (−0.005, +0.005)	−0.000 (−0.005, +0.005)	−0.000 (−0.006, +0.005)
Cancer	PFOS	−0.001 (−0.013, +0.010)	−0.001 (−0.011, +0.009)	−0.000 (−0.012, +0.012)	+0.001 (−0.009, +0.013)
Cancer	PFOA	+0.001 (−0.011, +0.012)	+0.003 (−0.006, +0.013)	+0.006 (−0.005, +0.017)	+0.004 (−0.004, +0.017)
Cancer	PFNA	−0.007 (−0.019, +0.005)	−0.006 (−0.016, +0.005)	−0.004 (−0.015, +0.008)	+0.003 (−0.005, +0.014)
Cancer	PFHxS	+0.000 (−0.008, +0.008)	−0.001 (−0.008, +0.007)	−0.001 (−0.010, +0.007)	−0.001 (−0.010, +0.005)
Cancer	PFDeA	+0.001 (−0.009, +0.011)	+0.001 (−0.008, +0.009)	+0.001 (−0.009, +0.010)	−0.000 (−0.009, +0.006)
Cancer	PFUnDA	−0.003 (−0.008, +0.002)	−0.003 (−0.008, +0.002)	−0.002 (−0.008, +0.003)	+0.001 (−0.004, +0.005)

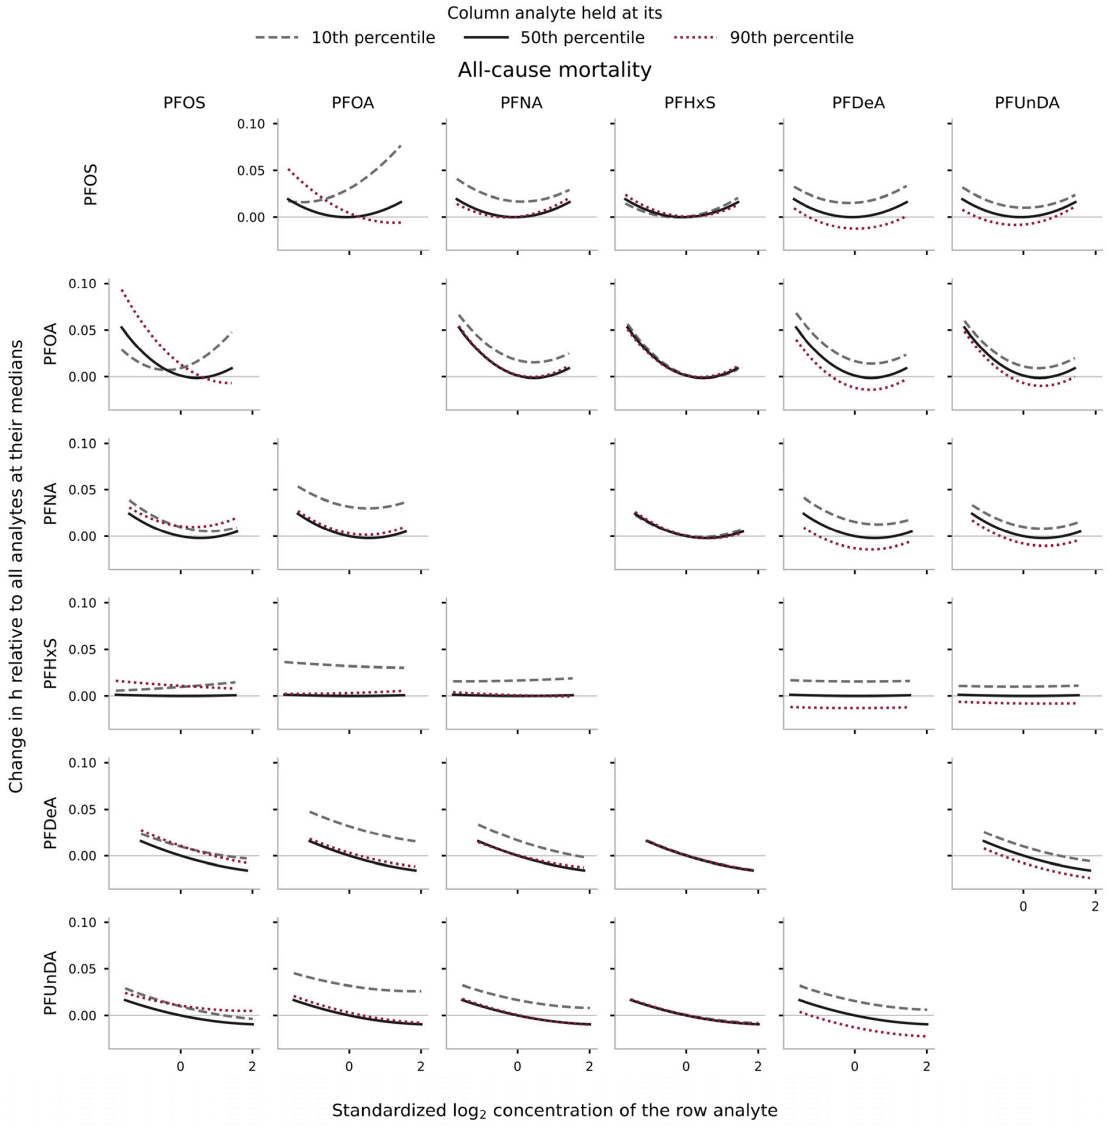

Figure 5. Bivariate exposure-response for all-cause mortality. The horizontal axis is the row analyte; the column analyte is held at its 10th, 50th or 90th percentile and the remaining four at their medians. Parallel curves indicate no interaction within that pair. Each point is referenced to the all-analytes-at-median point within the same posterior draw, which removes the unidentified additive constant in the surface. PFOS, perfluorooctane sulfonate; PFOA, perfluorooctanoic acid; PFNA, perfluorononanoic acid; PFHxS, perfluorohexane sulfonic acid; PFDeA, perfluorodecanoic acid; PFUnDA, perfluoroundecanoic acid. The cardiovascular and cancer surfaces are given as Figures S2 and S3 of the Supplementary Materials.

3.7. Shape of the Joint Exposure-Response on the Hazard Scale

The shape of the surface was confirmed by partial likelihood, without any approximation of the survival outcome. For all-cause mortality, natural cubic splines in the six analytes improved on the linear additive model (chi-squared 26.7 on 12 degrees of freedom, $p = 0.009$), while adding all fifteen pairwise products did not (chi-squared 17.3 on 15 degrees of freedom, $p = 0.299$). The joint exposure-response from the spline model declined monotonically: relative to all six analytes at their joint 50th percentile, the hazard ratio (HR) was 1.49 (1.29, 1.72) with all six at their joint 10th percentile, 1.19 (1.11, 1.27) at the 25th, 0.94 (0.89, 1.00) at the 75th and 0.85 (0.75, 0.96) at the 90th (Figure 4, Table 9). For cardiovascular mortality there was no departure from linearity (chi-squared 12.3 on 12 degrees of freedom, $p = 0.420$) and the pairwise products were borderline (chi-squared

24.6 on 15 degrees of freedom, $p = 0.056$); for cancer mortality neither test was significant (chi-squared 15.8 on 12 degrees of freedom, $p = 0.202$ and chi-squared 11.1 on 15 degrees of freedom, $p = 0.743$). The kernel machine surface and the spline hazard model therefore agree on both the direction and the shape of the joint association for all-cause mortality and on the absence of a consistent joint association for the two cause-specific endpoints, and the conclusion about shape does not rest on the Gaussian working approximation.

Table 9. Likelihood ratio tests from the spline Cox models, comparing a model linear and additive in the six base-two logarithmic concentrations with one allowing natural cubic splines in each analyte, and that model with one adding all fifteen pairwise products. CVD, cardiovascular disease.

Outcome	Test	LR chi2	df	*p*
All-cause	Nonlinearity (splines vs. linear additive)	26.69	12	0.009
All-cause	Interaction (pairwise products added)	17.34	15	0.299
CVD	Nonlinearity (splines vs. linear additive)	12.32	12	0.420
CVD	Interaction (pairwise products added)	24.58	15	0.056
Cancer	Nonlinearity (splines vs. linear additive)	15.76	12	0.202
Cancer	Interaction (pairwise products added)	11.14	15	0.743

4. Discussion

In this general-population sample, a higher joint serum PFAS (per- and polyfluoroalkyl substances) burden was associated with lower all-cause mortality, with no clear evidence of an association with cardiovascular or cancer mortality. The all-cause association was consistent across the specifications examined. It held on the hazard scale with individual follow-up time; under adjustment for or omission of survey cycle; under adjustment for smoking status, the family poverty–income ratio, alcohol use, physical activity, baseline disease, seafood consumption and energy intake; with serum cotinine in place of self-reported smoking; with the PFAS subsample weights and the variance clustered on the primary sampling unit; under landmarks at two and five years; when the mixture was restricted to the four analytes with near-complete detection; and under both a flexible kernel machine surface and a spline hazard model, which placed the joint exposure-response credibly below the null at high exposure.

The quantity being estimated deserves to be stated plainly, because it is not a single-compound effect. The joint estimate is the hazard ratio comparing the hazard that would be observed if every participant's exposure were shifted up by one quartile on all six analytes at once against the hazard under their observed exposure. It is a single number for a simultaneous move of the whole mixture, and reading it requires three things to hold. Shifting the whole mixture must be a contrast the data support, which the strong positive correlation among the analytes makes reasonable here, since high and low values of the analytes co-occur in the same people. The model must be correctly specified for the joint movement, which the spline tests support for a monotone surface in these data. And the analyte weights must be read as a description of how the joint estimate decomposes within this model, not as independent effects: they are shares of a signed partial sum and they change if the set of analytes changes. The estimate answers what happens when the mixture moves as a whole, which is how these compounds actually occur, and it does not answer what happens when one of them moves alone.

This should not be read as protection. Three explanations are more plausible than a protective effect of exposure, and the analyses allow them to be ranked by the support the data give them. None is established by these data.

The first is reverse causation. The association was numerically larger in the first decade of follow-up, 0.80 (0.73, 0.88), than afterwards, 0.96 (0.80, 1.16), and it attenuated from 0.83 to 0.88 under a five-year landmark. That pattern is the one reverse causation predicts. The evidence for it is limited, however. The difference between the two periods was not statistically significant ($p = 0.070$), and the association remained significant under the five-year landmark. Serum PFAS concentrations can be altered by processes that precede death from chronic disease, including reduced dietary intake, weight loss and falling serum albumin, none of which reflect exposure history.

The second, and the one these data support most directly, is that serum PFAS concentration is itself partly determined by physiological status. These compounds bind serum albumin and are cleared by the kidney, so albumin concentration and glomerular filtration both enter the measured exposure. Adjustment for serum albumin and the estimated glomerular filtration rate (eGFR) attenuated the all-cause estimate from 0.83 to 0.87, about a quarter of the association on the logarithmic scale, and it was the only specification among those examined that moved the estimate materially. Albumin and eGFR are also markers of the subclinical illness that precedes death, so this attenuation and the time structure in the preceding paragraph are two views of the same mechanism rather than independent findings. It also means the adjustment cannot be interpreted cleanly, because conditioning on a variable that lies between exposure measurement and death removes some genuine exposure signal along with the artifact.

The third is socioeconomic confounding, and here the result is informative in the opposite direction to what was expected. The crude configuration is present and strong: across household income brackets the mixture score rose and mortality fell steeply, which is exactly the structure that generates an inverse exposure-outcome association in the absence of any causal relation. Adjustment, however, did not attenuate the estimate. Adding household income left it at 0.83 against 0.82 without it, and the poverty–income ratio moved the primary estimate from 0.82 to 0.83, slightly away from the null. The dietary pathway behaves the same way. Seafood consumption tracks both higher exposure and higher socioeconomic position, with the mean mixture quartile score rising from 7.32 to 9.92 and the mean poverty–income ratio from 2.05 to 2.95 across categories of fish and shellfish meals, yet adjusting for seafood consumption and energy intake left the estimate at 0.79 against 0.80 in the same participants. Income, the poverty–income ratio and a 30-day seafood frequency are still coarse measures of the socioeconomic gradient and of habitual diet, so this does not exclude confounding operating through unmeasured pathways. It does mean that the socioeconomic explanation, which is the one most often offered for inverse PFAS associations in general-population samples, is not supported by the measures available in these data.

PFAS were measured once at baseline, and a single measurement may not represent exposure over the full follow-up. The long-chain compounds studied here have long serum elimination half-lives, approximately 2.7 to 3.5 years for PFOA (perfluorooctanoic acid), 3.4 to 5.7 years for PFOS (perfluorooctane sulfonate) and 5.3 years for PFHxS (perfluorohexane sulfonic acid) [6,7], with fewer data for PFNA (perfluorononanoic acid), PFDeA (perfluorodecanoic acid) and PFUnDA (perfluoroundecanoic acid). A single measurement therefore reflects exposure accumulated over the preceding years reasonably well, and is less informative about exposure later in follow-up. Population concentrations changed during the study period: across the four cycles analyzed here the geometric mean of PFOS fell from 20.04 to 9.33 ng/mL while PFNA rose from 0.91 to 1.24 ng/mL (Table 1). A participant's rank at baseline may therefore drift over time, and because the half-lives differ between compounds the drift is likely to differ between them as well, so the composition of the mixture a participant carries can change even when their rank does not. Such misclassification would

tend to grow with follow-up time and may contribute to the weaker association observed after ten years.

The absence of a clear cardiovascular association sits against single-compound reports of positive PFOS and PFOA associations with cardiovascular mortality [28,29]. Part of the difference is the mixture framing: in a joint model the correlated analytes compete, and the positive and negative partial contributions to the cardiovascular estimate here were of similar size. Part is exposure range and exposure history. General-population concentrations in the National Health and Nutrition Examination Survey are far below those in the contaminated-area cohorts where positive mortality associations have been observed, and those studies assign exposure by residence over decades rather than by a single serum measurement [28,29]. The two designs address different exposure histories and should not be read as replications of one another. It is also worth noting that the per-doubling cardiovascular estimate, 0.88 (0.76, 1.01), and the PFOA-specific estimate, 0.82 (0.69, 0.96), point in the same inverse direction as the all-cause result, which is more consistent with the shared non-causal explanations above than with a cardiovascular-specific effect.

The analyte pattern should be separated from the mixture result. PFOA (perfluorooctanoic acid) was the only analyte with both individual and posterior support, and it carried the largest share of the negative partial sum in the quantile g-computation decomposition. Posterior inclusion probabilities describe selection support within the fitted surface and are not effect sizes, and the weights are relative shares of a partial sum rather than independent estimates. Neither licenses a ranking of these compounds by toxicity. The point carries particular force here because the analytes are correlated: with Spearman coefficients up to 0.75, component-wise selection tends to concentrate support on one representative of a correlated set, so a high posterior inclusion probability for PFOA and near-zero values for the others is evidence about which analyte best carries the common signal, not evidence that the others are inert.

A methodological point generalizes beyond this dataset. Pooling cycles of the National Health and Nutrition Examination Survey for a mortality endpoint gives participants systematically different lengths of observation, and exposures that change over calendar time are therefore entangled with the window in which deaths are counted. Using individual follow-up time in a survival model resolves this directly, which is why cycle adjustment changed the estimates here very little. Analyses that instead model a binary mortality indicator over cycle-specific windows remain vulnerable to it, particularly for mixtures whose components all trend in one direction over calendar time, and the instability of the probit specification reported in Section 3.5 is one symptom of discarding that information.

Strengths and Limitations

This study has several methodological strengths. Mortality follow-up was complete for every participant, so all deaths were analyzed on the hazard scale with individual time at risk. Covariates covering smoking, socioeconomic position, alcohol use, physical activity, baseline disease and two laboratory markers of baseline physiological status were available for almost the whole sample, and design-based analyses used the PFAS (per- and polyfluoroalkyl substances) subsample weights with the variance clustered on the primary sampling unit. The time structure of the association was examined through landmark and period-specific analyses with correctly constructed risk sets. The mixture method imposes no directional constraint. Competing-risks models were fitted for the cause-specific endpoints. The shape of the joint exposure-response was established on the hazard scale by partial likelihood as well as by a flexible kernel machine surface fitted to the

full cohort with multi-chain convergence diagnostics for every reported summary, and the two agreed.

Several limitations qualify the findings. Serum PFAS were measured once, so within-person change over follow-up is unmeasured; non-differential measurement error of this kind would attenuate estimates, and its magnitude is likely to grow with follow-up time. Dietary intake was characterized by the number of fish and shellfish meals in the past 30 days together with total energy intake, which captures the pathway through which residual socioeconomic confounding is most likely to operate, since higher PFAS concentrations track consumption of seafood and commercially prepared food, which in turn tracks income and education; a 30-day frequency is nonetheless a coarse measure of habitual intake, and other dietary sources, including food packaging, were not measured. Occupational and other environmental exposures were not measured either; serum concentrations integrate exposure from all routes, so these sources are captured in the exposure measure itself, but their independent contributions to mortality are not. Adjustment for serum albumin and the estimated glomerular filtration rate addresses one mechanism by which illness alters measured exposure, but both lie between exposure measurement and death, so that adjustment cannot be given a clean causal reading. The primary analysis excluded the 467 participants with an incomplete covariate record, almost all of them missing the poverty–income ratio, and the full-sample estimate under the reduced covariate set is reported alongside for comparison. Design-based estimates use the correct subsample weights, but the analytic sample is a one-third subsample that is healthier than the full survey population, so selection into it may contribute. Cause of death is taken from death certificates and is subject to misclassification, which would bias the cause-specific estimates toward the null, and with 344 cardiovascular and 235 cancer deaths the cause-specific analyses have limited power: the cardiovascular confidence intervals admit hazard ratios from about 0.79 to 1.07 per quartile increase, so a modest association in either direction cannot be excluded. Finally, the three mortality endpoints and two exposure scalings were examined without adjustment for multiplicity, as were the eighteen interaction summaries, so p-values and intervals close to the null should be read as descriptive.

5. Conclusions

A higher joint serum PFAS (per- and polyfluoroalkyl substances) burden was associated with lower all-cause mortality in this general-population cohort, consistently across survival, spline and kernel machine approaches, and there was no clear evidence of an association with either cardiovascular or cancer mortality. The inverse all-cause association should not be interpreted as evidence that PFAS exposure is protective. It was larger in the first decade of follow-up than afterwards and attenuated under lag exclusion, and it attenuated further on adjustment for serum albumin and the estimated glomerular filtration rate, which are determinants of the measured exposure as well as markers of baseline health. These patterns are consistent with reverse causation and with physiological determinants of serum concentration contributing to the inverse direction, although neither is established by these data. Measured socioeconomic position had the structure that could generate such an association but did not account for it, so the common socioeconomic explanation is not supported here. Quantifying what survives these explanations will require repeated exposure measurement and measures of diet that were not available. Separately, these data show that pooled-cycle mortality analyses in the National Health and Nutrition Examination Survey should use individual follow-up time rather than a binary outcome over cycle-specific windows, because exposures that trend over calendar time are otherwise entangled with the window in which deaths are counted.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/jox16080229/s1. Table S1: Convergence diagnostics for every Bayesian kernel machine regression model and every reported parameter. Figure S1: Bayesian kernel machine regression single-analyte effects, from the fits with every analyte forced into the surface. Figure S2: Bivariate exposure-response for cardiovascular mortality. Figure S3: Bivariate exposure-response for cancer mortality.

Author Contributions: Conceptualization, E.O.-G.; methodology, E.O. and E.O.-G.; formal analysis, E.O. and E.O.-G.; investigation, E.O. and E.O.-G.; resources, E.O.-G.; data curation, E.O. and E.O.-G.; writing—original draft preparation, E.O. and E.O.-G.; writing—review and editing, E.O. and E.O.-G.; supervision, E.O.-G.; project administration, E.O.-G.; funding acquisition, E.O.-G. All authors have read and agreed to the published version of the manuscript.

Funding: Research reported in this publication was supported by the National Institute of General Medical Sciences of the National Institutes of Health under award number R16GM149473. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

Institutional Review Board Statement: Not applicable. This secondary analysis used de-identified public-use data and did not require additional review.

Informed Consent Statement: Not applicable. This study is a secondary analysis of public data and does not involve new human research. No additional informed consent was required.

Data Availability Statement: All data analyzed are publicly available. Examination, laboratory and questionnaire files of the National Health and Nutrition Examination Survey are at https://wwwn.cdc.gov/nchs/nhanes/ (accessed on 26 October 2026), and the public-use linked mortality files of the National Center for Health Statistics at https://www.cdc.gov/nchs/linked-data/mortality-files/index.html (accessed on 26 October 2026).

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

Abbreviation	Definition
PFAS	per- and polyfluoroalkyl substances
PFOA	perfluorooctanoic acid
PFOS	perfluorooctane sulfonate
PFNA	perfluorononanoic acid
PFHxS	perfluorohexane sulfonic acid
PFDeA	perfluorodecanoic acid
PFUnDA	perfluoroundecanoic acid
NHANES	National Health and Nutrition Examination Survey
NCHS	National Center for Health Statistics
BKMR	Bayesian kernel machine regression
BMI	body mass index
CI	confidence interval
CrI	credible interval
CVD	cardiovascular disease
eGFR	estimated glomerular filtration rate
GM	geometric mean
HR	hazard ratio
IARC	International Agency for Research on Cancer
IQR	interquartile range
LOD	limit of detection

PIP	posterior inclusion probability
PIR	poverty–income ratio
PSU	primary sampling unit

References

1. Lindstrom, A.B.; Strynar, M.J.; Libelo, E.L. Polyfluorinated compounds: Past, present, and future. *Environ. Sci. Technol.* **2011**, *45*, 7954–7961. [CrossRef] [PubMed]
2. Wang, Z.; DeWitt, J.C.; Higgins, C.P.; Cousins, I.T. A never-ending story of per- and polyfluoroalkyl substances (PFASs)? *Environ. Sci. Technol.* **2017**, *51*, 2508–2518. [CrossRef] [PubMed]
3. Buck, R.C.; Franklin, J.; Berger, U.; Conder, J.M.; Cousins, I.T.; de Voogt, P.; Jensen, A.A.; Kannan, K.; Mabury, S.A.; van Leeuwen, S.P. Perfluoroalkyl and polyfluoroalkyl substances in the environment: Terminology, classification, and origins. *Integr. Environ. Assess. Manag.* **2011**, *7*, 513–541. [CrossRef] [PubMed]
4. Sunderland, E.M.; Hu, X.C.; Dassuncao, C.; Tokranov, A.K.; Wagner, C.C.; Allen, J.G. A review of the pathways of human exposure to poly- and perfluoroalkyl substances (PFASs) and present understanding of health effects. *J. Expo. Sci. Environ. Epidemiol.* **2019**, *29*, 131–147. [CrossRef] [PubMed]
5. Ng, C.A.; Hungerbühler, K. Bioconcentration of perfluorinated alkyl acids: How important is specific binding? *Environ. Sci. Technol.* **2013**, *47*, 7214–7223. [CrossRef] [PubMed]
6. Rosato, I.; Bonato, T.; Fletcher, T.; Batzella, E.; Canova, C. Estimation of per- and polyfluoroalkyl substances (PFAS) half-lives in human studies: A systematic review and meta-analysis. *Environ. Res.* **2024**, *242*, 117743. [CrossRef] [PubMed]
7. Li, Y.; Fletcher, T.; Mucs, D.; Scott, K.; Lindh, C.H.; Tallving, P.; Jakobsson, K. Half-lives of PFOS, PFHxS and PFOA after end of exposure to contaminated drinking water. *Occup. Environ. Med.* **2018**, *75*, 46–51. [CrossRef] [PubMed]
8. Jain, R.B. Time trends over 2003–2014 in the concentrations of selected perfluoroalkyl substances among US adults aged ≥20 years: Interpretational issues. *Sci. Total Environ.* **2018**, *645*, 946–957. [CrossRef] [PubMed]
9. Calafat, A.M.; Wong, L.Y.; Kuklenyik, Z.; Reidy, J.A.; Needham, L.L. Polyfluoroalkyl chemicals in the U.S. population: Data from the National Health and Nutrition Examination Survey (NHANES) 2003–2004 and comparisons with NHANES 1999–2000. *Environ. Health Perspect.* **2007**, *115*, 1596–1602. [CrossRef] [PubMed]
10. United Nations Environment Programme. *Stockholm Convention on Persistent Organic Pollutants. Annex B (Restriction): PFOS*; United Nations Environment Programme: Geneva, Switzerland, 2009; updated 2019 (PFOA).
11. Domingo, J.L.; Nadal, M. Human exposure to per- and polyfluoroalkyl substances (PFAS) through drinking water: A review of the recent scientific literature. *Environ. Res.* **2019**, *177*, 108648. [CrossRef] [PubMed]
12. Hu, X.C.; Andrews, D.Q.; Lindstrom, A.B.; Bruton, T.A.; Schaider, L.A.; Grandjean, P.; Lohmann, R.; Carignan, C.C.; Blum, A.; Balan, S.A.; et al. Detection of poly- and perfluoroalkyl substances (PFASs) in U.S. drinking water linked to industrial sites, military fire training areas, and wastewater treatment plants. *Environ. Sci. Technol. Lett.* **2016**, *3*, 344–350. [CrossRef] [PubMed]
13. Domingo, J.L.; Nadal, M. Per- and polyfluoroalkyl substances (PFASs) in food and human dietary intake: A review of the recent scientific literature. *J. Agric. Food Chem.* **2017**, *65*, 533–543. [CrossRef] [PubMed]
14. Noorlander, C.W.; van Leeuwen, S.P.; te Biesebeek, J.D.; Mengelers, M.J.; Zeilmaker, M.J. Levels of perfluorinated compounds in food and dietary intake of PFOS and PFOA in the Netherlands. *J. Agric. Food Chem.* **2011**, *59*, 7496–7505. [CrossRef] [PubMed]
15. Fragki, S.; Dirven, H.; Fletcher, T.; Grasl-Kraupp, B.; Bjerve Gützkow, K.; Hoogenboom, R.; Kersten, S.; Lindeman, B.; Louisse, J.; Peijnenburg, A.; et al. Systemic PFOS and PFOA exposure and disturbed lipid homeostasis in humans: What do we know and what not? *Crit. Rev. Toxicol.* **2021**, *51*, 141–164. [CrossRef] [PubMed]
16. Rosen, M.B.; Abbott, B.D.; Wolf, D.C.; Corton, J.C.; Wood, C.R.; Schmid, J.E.; Das, K.P.; Zehr, R.D.; Blair, E.T.; Lau, C. Gene profiling in the livers of wild-type and PPARα-null mice exposed to perfluorooctanoic acid. *Toxicol. Pathol.* **2008**, *36*, 592–607. [CrossRef] [PubMed]
17. Kowalska, D.; Sosnowska, A.; Bulawska, N.; Stępnik, M.; Besselink, H.; Behnisch, P.; Puzyn, T. How the structure of per- and polyfluoroalkyl substances (PFAS) influences their binding potency to the peroxisome proliferator-activated and thyroid hormone receptors: An in silico screening study. *Molecules* **2023**, *28*, 479. [CrossRef] [PubMed]
18. National Academies of Sciences, Engineering, and Medicine. *Guidance on PFAS Exposure, Testing, and Clinical Follow-Up*; National Academies Press: Washington, DC, USA, 2022. [CrossRef] [PubMed]
19. Grandjean, P.; Timmermann, C.A.G.; Kruse, M.; Nielsen, F.; Vinholt, P.J.; Boding, L.; Heilmann, C.; Mølbak, K. Severity of COVID-19 at elevated exposure to perfluorinated alkylates. *PLoS ONE* **2020**, *15*, e0244815. [CrossRef] [PubMed]
20. Ji, C.; Pavuk, M.; Ruiz, P. Global serum per- and polyfluoroalkyl substances exposures and their correlation with lipids: A systematic review and meta-analysis. *Environ. Sci. Technol.* **2025**, *59*, 19611–19629. [CrossRef] [PubMed]
21. Christensen, B.T.; Calkins, M.M. Occupational exposure to per- and polyfluoroalkyl substances: A scope review of the literature from 1980–2021. *J. Expo. Sci. Environ. Epidemiol.* **2023**, *33*, 673–686. [CrossRef] [PubMed]

22. Liu, G.; Dhana, K.; Furtado, J.D.; Rood, J.; Zong, G.; Liang, L.; Qi, L.; Bray, G.A.; DeJonge, L.; Coull, B.; et al. Perfluoroalkyl substances and changes in body weight and resting metabolic rate in response to weight-loss diets: A prospective study. *PLoS Med*. **2018**, *15*, e1002502. [CrossRef] [PubMed]
23. Costello, E.; Rock, S.; Stratakis, N.; Eckel, S.P.; Walker, D.I.; Valvi, D.; Cserbik, D.; Jenkins, T.; Xanthakos, S.A.; Kohli, R.; et al. Exposure to per- and polyfluoroalkyl substances and markers of liver injury: A systematic review and meta-analysis. *Environ. Health Perspect*. **2022**, *130*, 46001. [CrossRef] [PubMed]
24. Pahwa, R.; Jialal, I. Atherosclerosis. In *StatPearls*; StatPearls Publishing: Treasure Island, FL, USA, 2024.
25. Higashi, Y. Roles of oxidative stress and inflammation in vascular endothelial dysfunction-related disease. *Antioxidants* **2022**, *11*, 1958. [CrossRef] [PubMed]
26. Zahm, S.; Bonde, J.P.; Chiu, W.A.; Hoppin, J.; Kanno, J.; Abdallah, M.; Blystone, C.R.; Calkins, M.M.; Dong, G.H.; Dorman, D.C.; et al. Carcinogenicity of perfluorooctanoic acid and perfluorooctanesulfonic acid. *Lancet Oncol*. **2024**, *25*, 16–17. [CrossRef] [PubMed]
27. International Agency for Research on Cancer. *Perfluorooctanoic Acid (PFOA) and Perfluorooctanesulfonic Acid (PFOS)*; IARC Monographs on the Identification of Carcinogenic Hazards to Humans; IARC: Lyon, France, 2025; Volume 135, pp. 1–754. Available online: https://publications.iarc.who.int/636 (accessed on 26 October 2026).
28. Biggeri, A.; Stoppa, G.; Facciolo, L.; Fin, G.; Mancini, S.; Manno, V.; Minelli, G.; Zamagni, F.; Zamboni, M.; Catelan, D.; et al. All-cause, cardiovascular disease and cancer mortality in the population of a large Italian area contaminated by perfluoroalkyl and polyfluoroalkyl substances (1980–2018). *Environ. Health* **2024**, *23*, 42. [CrossRef] [PubMed]
29. Zhou, L.; Andersson, E.M.; Harari, F.; Fletcher, T.; Nielsen, C.; Blomberg, A.J.; Jakobsson, K.; Xu, Y.; Li, Y. Exposure to high levels of perfluoroalkyl substances through drinking water and risk of cardiovascular morbidity and mortality in a Swedish register-based study. *Environ. Res*. **2025**, *286*, 122765. [CrossRef] [PubMed]
30. Dunder, L.; Salihovic, S.; Varotsis, G.; Lind, P.M.; Elmståhl, S.; Lind, L. Plasma levels of per- and polyfluoroalkyl substances (PFAS) and cardiovascular disease: Results from two independent population-based cohorts and a meta-analysis. *Environ. Int*. **2023**, *181*, 108250. [CrossRef] [PubMed]
31. Steenland, K.; Winquist, A. PFAS and cancer, a scoping review of the epidemiologic evidence. *Environ. Res*. **2021**, *194*, 110690. [CrossRef] [PubMed]
32. Winquist, A.; Hodge, J.M.; Diver, W.R.; Rodriguez, J.L.; Troeschel, A.N.; Daniel, J.; Teras, L.R. Case-cohort study of the association between PFAS and selected cancers among participants in the American Cancer Society's Cancer Prevention Study II LifeLink cohort. *Environ. Health Perspect*. **2023**, *131*, 127007. [CrossRef] [PubMed]
33. Seyyedsalehi, M.S.; Boffetta, P. Per- and poly-fluoroalkyl substances (PFAS) exposure and risk of kidney, liver, and testicular cancers: A systematic review and meta-analysis. *Med. Lav*. **2023**, *114*, e2023040. [CrossRef] [PubMed]
34. Zheng, J.; Liu, S.; Yang, J.; Zheng, S.; Sun, B. Per- and polyfluoroalkyl substances (PFAS) and cancer: Detection methodologies, epidemiological insights, potential carcinogenic mechanisms, and future perspectives. *Sci. Total Environ*. **2024**, *953*, 176158. [CrossRef] [PubMed]
35. Braun, J.M.; Gennings, C.; Hauser, R.; Webster, T.F. What can epidemiological studies tell us about the impact of chemical mixtures on human health? *Environ. Health Perspect*. **2016**, *124*, A6–A9. [CrossRef] [PubMed]
36. Johnson, C.L.; Paulose-Ram, R.; Ogden, C.L.; Carroll, M.D.; Kruszon-Moran, D.; Dohrmann, S.M.; Curtin, L.R. National Health and Nutrition Examination Survey: Analytic guidelines, 1999–2010. *Vital Health Stat. 2* **2013**, *161*, 1–24.
37. National Center for Health Statistics. The NCHS Public-Use Linked Mortality Files. Available online: https://www.cdc.gov/nchs/linked-data/mortality-files/index.html (accessed on 26 October 2026).
38. Keil, A.P.; Buckley, J.P.; O'Brien, K.M.; Ferguson, K.K.; Zhao, S.; White, A.J. A quantile-based g-computation approach to addressing the effects of exposure mixtures. *Environ. Health Perspect*. **2020**, *128*, 047004. [CrossRef] [PubMed]
39. Carrico, C.; Gennings, C.; Wheeler, D.C.; Factor-Litvak, P. Characterization of weighted quantile sum regression for highly correlated data in a risk analysis setting. *J. Agric. Biol. Environ. Stat*. **2015**, *20*, 100–120. [CrossRef] [PubMed]
40. National Center for Health Statistics. Analytic Guidelines for the National Health and Nutrition Examination Survey. Available online: https://wwwn.cdc.gov/nchs/nhanes/analyticguidelines.aspx (accessed on 26 October 2026).
41. Bobb, J.F.; Valeri, L.; Claus Henn, B.; Christiani, D.C.; Wright, R.O.; Mazumdar, M.; Godleski, J.J.; Coull, B.A. Bayesian kernel machine regression for estimating the health effects of multi-pollutant mixtures. *Biostatistics* **2015**, *16*, 493–508. [CrossRef] [PubMed]
42. Bobb, J.F.; Claus Henn, B.; Valeri, L.; Coull, B.A. Statistical software for analyzing the health effects of multiple concurrent exposures via Bayesian kernel machine regression. *Environ. Health* **2018**, *17*, 67. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.