

inedemogR: Tidy Access to Spanish INE Demographic Data

A Comprehensive Tutorial

J. R. Caro-Barrera

August 20, 2026

Contents

1	Introduction	3
1.1	Installation	3
1.2	A note on live data	3
2	Part I — Data Retrieval, Cleaning, and Storage	5
2.1	System A: quick multi-geography totals	5
2.1.1	list_ine_indicators()	5
2.1.2	get_ine_demog()	5
2.1.3	get_ine_geo()	6
2.1.4	update_ine_data()	6
2.2	System B: age/sex-disaggregated province data	6
2.2.1	Caching layer (helpers.R)	6
2.2.2	get_ine_population()	7
2.2.3	get_ine_births()	7
2.2.4	get_ine_births_by_age()	7
2.2.5	get_ine_deaths()	8
2.3	The mortality pipeline	8
2.3.1	Exposure-to-risk: compute_exposure()	8
2.3.2	Central death rates: compute_death_rates()	9
2.3.3	Period life tables: build_life_tables()	10
2.3.4	Abridged life tables: build_abridged_life_tables()	11
2.4	Exporting to files: download_ine_data()	12
3	Part II — Demographic Analysis and Visualization	13
3.1	Summary demographic indicators	13
3.1.1	Age dependency ratios: age_dependency_ratio()	13
3.1.2	Aging index: aging_index()	13
3.1.3	Sex ratio: sex_ratio()	13
3.1.4	Crude birth and death rates	13
3.1.5	General fertility rate: general_fertility_rate()	14
3.1.6	Fertility schedule: ASFR, TFR, MAC, GRR, NRR	14
3.1.7	Total fertility rate: total_fertility_rate()	14
3.1.8	Mean age at childbearing: mean_age_at_childbearing()	15

3.1.9	Gross and net reproduction rates	15
3.1.10	Infant mortality rate: <code>infant_mortality_rate()</code>	15
3.1.11	Rate of natural increase: <code>rate_of_natural_increase()</code>	16
3.1.12	Birth-to-death ratio: <code>birth_death_ratio()</code>	16
3.1.13	Life expectancy: <code>life_expectancy_summary()</code> and <code>life_expectancy()</code>	16
3.2	Mortality decomposition: <code>decompose_life_expectancy()</code>	16
3.3	Mapping functions	18
3.3.1	<code>plot_in_e_map()</code>	18
3.3.2	<code>map_indicator()</code>	18
3.3.3	<code>map_life_expectancy()</code>	18
3.4	Population pyramids and trend charts	18
3.4.1	<code>plot_population_pyramid()</code>	18
3.4.2	<code>plot_demog_trend()</code>	19
3.5	Lexis diagrams: <code>plot_lexis_diagram()</code>	19
3.6	Worked examples	19
3.6.1	Birth-to-death ratio map (<code>birth_death_ratio_map.R</code>)	19
3.6.2	Andalusian population pyramids (<code>andalusia_pyramids.R</code>)	20
3.6.3	Madrid/Barcelona Lexis diagrams (<code>madrid_barcelona_lexis.R</code>)	20
3.6.4	Life expectancy comparison (<code>life_expectancy_comparison.R</code>)	21
3.6.5	Natural increase analysis (<code>natural_increase_analysis.R</code>)	21
3.6.6	Fertility and infant mortality analysis (<code>fertility_infant_mortality_analysis.R</code>)	21
3.6.7	Fertility indicators analysis (<code>fertility_indicators_analysis.R</code>)	22
3.6.8	Mortality decomposition analysis (<code>mortality_decomposition_analysis.R</code>)	23
3.6.9	Abridged life tables analysis (<code>abridged_life_tables_analysis.R</code>)	23

Appendix

25

1 Introduction

inedemogR provides tidy, harmonized, reproducible access to demographic data published by the Spanish National Statistics Institute (Instituto Nacional de Estadística, INE), retrieved live via the official **ineapir** API wrapper, with spatial integration via **mapSpain** and plotting built on **ggplot2**.

The package is organized around two complementary data systems that serve different purposes. They are deliberately kept parallel rather than merged into one interface, because they operate at genuinely different levels of demographic detail.

System A `get_ine_demog()`, `get_ine_geo()`, `list_ine_indicators()`, `plot_ine_map()`, `map_indicator()`, `update_ine_data()`. Quick, multi-geography (municipality or province) totals for population, births, and deaths, with no age or sex breakdown. Best for national/regional choropleths and quick comparisons.

System B `get_ine_population()`, `get_ine_births()`, `get_ine_births_by_age()`, `get_ine_deaths()`, `compute_exposure()`, `compute_death_rates()`, `build_life_tables()`, `build_abridged_life_tables()`, `download_ine_data()`, and the summary indicators, decomposition, and charts built on top of them. Rigorous, province-level, age/sex-disaggregated analysis following the Human Mortality Database (HMD) Methods Protocol V6, as implemented for the Spanish subnational Human Mortality Database (SHMD) pipeline. This is the only system capable of producing dependency ratios, population pyramids, Lexis diagrams, period life tables, fertility schedules (ASFR/TFR/MAC/GRR/NRR), and mortality decomposition, since these all require an age breakdown that System A's data does not carry.

This tutorial is split into two parts, matching the two purposes a user of **inedemogR** typically has:

- **Part I (Section 2)** covers strictly data *retrieval, cleaning, validation, and storage* — the functions in `get_ine_demog.R`, `get_ine_geo.R`, `population.R`, `births.R`, `deaths.R`, `exposure.R`, `death_rates.R`, `life_tables.R`, `abridged_life_tables.R`, `download_ine_data.R`, and the caching layer in `helpers.R`.
- **Part II (Section 3)** covers *demographic analysis and visualization* — the indicator functions in `indicators.R`, the mortality-decomposition functions in `decomposition.R`, and the charting/mapping functions in `plot_demog.R`, `plot_ine_map.R`, and `lexis.R` — together with full worked examples.

Every mortality-pipeline calculation (exposure-to-risk, central death rates, life tables) is presented with the exact mathematical formulation implemented in the code, following the HMD Methods Protocol V6.

1.1 Installation

```
# From a local clone or source tarball:
# install.packages("devtools")
devtools::install()

library(inedemogR)
```

1.2 A note on live data

Every retrieval function in this package calls the real INE API (or, for `get_ine_geo()`, the real **mapSpain** boundary service) over the network. There is no bundled offline dataset. All code listings in this tutorial are real, runnable calls against live INE data as of the 2023–2024 reference years used in the examples;

results will naturally update as INE publishes new years. Section [2.2.1](#) explains how to avoid repeatedly re-downloading the same data during iterative analysis.

2 Part I — Data Retrieval, Cleaning, and Storage

2.1 System A: quick multi-geography totals

2.1.1 `list_ine_indicators()`

Returns the internal registry (`ine_variables`) mapping each indicator code to the real INE table it is retrieved from, the geographic dimension column INE's API returns for that table, and the geographic granularity the table is actually published at.

```
list_ine_indicators()
#> # A tibble: 3 x 6
#>   indicator      name      description id_table geo_var geo_level
#>   population_total Resident pop. ...      29005 Municipios municipality
#>   births_total    Live births    ...      6506 Provincias province
#>   deaths_total    Deaths        ...      6545 Provincias province
```

Migration indicators are not included in this release: INE's migration tables are origin×destination×year flow matrices, a different shape from the simple geography×year stock tables backing the three indicators above, so wiring one in needs its own design pass rather than a registry one-liner. `get_ine_demog()` still raises an explicit "Unknown indicator(s)" error for any indicator code not in `list_ine_indicators()`, migration or otherwise.

2.1.2 `get_ine_demog()`

Retrieves one or more indicators via `ineapir::get_data_table()` and tidies the result into one row per geography × year, one column per indicator.

```
get_ine_demog(indicator, geo_level = NULL, year = NULL,
               region = NULL, sex = "Total", geometry = FALSE)
```

- | | |
|------------------|---|
| indicator | Character vector of indicator codes (see <code>list_ine_indicators()</code>). |
| geo_level | The geographic level to fetch; inferred from <code>indicator</code> if <code>NULL</code> . All requested indicators must share the same level (population is municipality-level; births/deaths are province-level only) — requesting a mix across levels is a hard error. |
| year | Numeric/character year; <code>NULL</code> returns only the latest published period. |
| region | Optional regex vector matched against the geography name. |
| sex | One of "Total", "Hombres", "Mujeres". |
| geometry | If <code>TRUE</code> , joins <code>get_ine_geo()</code> 's geometry and returns an <code>sf</code> object. |

```
# Population totals for every municipality, 2023
pop_totals <- get_ine_demog(indicator = "population_total", year = 2023)

# Births and deaths together, province level, one call
vitals <- get_ine_demog(
  indicator = c("births_total", "deaths_total"), year = 2023
)
```

2.1.3 get_ine_geo()

Retrieves province or municipality boundaries via **mapSpain** (`esp_get_prov()` / `esp_get_munic()`), re-projected to ETRS89 / UTM zone 30N (EPSG:25830), with a GEOID column that matches `get_ine_demog()`'s output exactly so the two can be joined directly.

```
get_ine_geo(geo_level = c("municipality", "province"), region = NULL,
            moveCAN = TRUE, can_gap_km = 60)
```

When `moveCAN = TRUE` (the default), the Canary Islands are translated so they sit `can_gap_km` kilometres from the nearest mainland coastline point (near Cádiz/Huelva) rather than left in their true, far-southwest position — the standard cartographic convention for compact national maps of Spain. The shifted islands' bounding box is attached as the `"can_box"` attribute, used by `plot_ine_map()` and `map_indicator()` to draw an inset separator frame.

```
prov_geo <- get_ine_geo(geo_level = "province")
mun_geo <- get_ine_geo(geo_level = "municipality", region = "Sevilla")
```

2.1.4 update_ine_data()

Checks every indicator wired in the registry for its latest published year and refreshes a local `.rds` cache (under `tools::R_user_dir("inedemogR", "cache")` by default) only if INE has published a newer year than what is cached, or if `force = TRUE`.

```
update_ine_data(force = FALSE, data_dir = tools::R_user_dir("inedemogR", "cache"))
```

2.2 System B: age/sex-disaggregated province data

2.2.1 Caching layer (helpers.R)

`get_ine_population()`, `get_ine_births()`, and `get_ine_deaths()` share a common disk-cache wrapper, `with_ine_cache()`. Fetching province-level age-specific deaths and population from INE's Tempus3 API (sometimes via a full CSV bulk-export fallback when the JSON API rejects an oversized query) is the slowest part of the package, often taking minutes; results are cached as `.rds` files keyed by the call's arguments so repeated calls during iterative analysis are near-instant.

Every System B retrieval function exposes the same three arguments:

- use_cache** (default `TRUE`) Return a previously cached result for the same arguments instead of re-fetching.
- force** (default `FALSE`) Ignore any existing cache and re-fetch from INE, refreshing the cache afterwards.
- cache_dir** Directory for the cache (default: `tools::R_user_dir("inedemogR", "cache")`, the same location `update_ine_data()` uses).

```
# First call: fetches live and caches the result.
pop <- get_ine_population(n_periods = 10)

# Second call, same arguments: returns instantly from cache.
pop <- get_ine_population(n_periods = 10)

# Force a fresh fetch, e.g. after INE has published a new year.
```

```
pop <- get_ine_population(n_periods = 10, force = TRUE)
```

2.2.2 get_ine_population()

Retrieves Spanish January-1st provincial population by single-year age (0–100+) and sex from INE's Padrón Municipal Continuo, per SHMD Protocol Section 12, Step 3.

```
get_ine_population(table_id = 56945, n_periods = 30,  
                  use_cache = TRUE, force = FALSE,  
                  cache_dir = tools::R_user_dir("inedemogR", "cache"))
```

Returns `list(data, qc)`: `data` has columns `ine_code`, `nuts3_code`, `nuts2_code`, `province_name`, `year`, `age`, `female`, `male`, `total`; `qc` is the output of `validate_population()`, which checks for negative/missing counts, incomplete age ranges, duplicate rows, year gaps, and implausible age-to-age jumps (>50%, informational only).

```
pop <- get_ine_population()  
head(pop$data)  
pop$qc$passed
```

2.2.3 get_ine_births()

Retrieves annual live births by province and sex from INE's MNPN operation.

```
get_ine_births(table_id = 6506, n_periods = 100,  
              use_cache = TRUE, force = FALSE,  
              cache_dir = tools::R_user_dir("inedemogR", "cache"))
```

Returns `list(data, qc)`; `data` has `ine_code`, `nuts3_code`, `nuts2_code`, `province_name`, `year`, `female`, `male`, `total`. `validate_births()` checks the same invariants as `validate_population()` (no age dimension here, so no age-range check).

2.2.4 get_ine_births_by_age()

`get_ine_births()` has no age-of-mother breakdown, so it cannot feed a true total fertility rate. `get_ine_births_by_age()` retrieves the same MNPN operation's age-of-mother table instead (table 6508: province, single year of age of the mother 15–49 plus the two open intervals under15/50plus, and sex of the newborn) — the input Section 3.1.6's age-specific fertility rate schedule needs.

```
get_ine_births_by_age(table_id = 6508, n_periods = 100,  
                    use_cache = TRUE, force = FALSE,  
                    cache_dir = tools::R_user_dir("inedemogR", "cache"))
```

Returns `list(data, qc)`; `data` has `ine_code`, `nuts3_code`, `nuts2_code`, `province_name`, `year`, `age` (integer, NA for the under15 group), `age_group` ("under15", "15" — "49", "50plus"), `female`, `male`, `total`. INE's table omits a (province, age, sex-of-newborn, year) series entirely once its value is zero across the whole fetched window — the same behavior already seen in `compute_death_rates()` (Section 2.3.2) — so `clean_births_by_age()` builds the complete grid explicitly and fills absent combinations with zero rather than leaving them missing, before `validate_births_by_age()` checks for negative/missing counts, duplicate rows, and complete single-year age coverage (15–49) per province-year.

```
births_age <- get_ine_births_by_age()
births_age$data
```

2.2.5 get_ine_deaths()

Retrieves annual deaths by province and sex (INE’s MNPD operation, table 6545) *and* an age-specific companion series (table 6547), since the mortality pipeline (`compute_exposure()`, `compute_death_rates()`) needs deaths by single-year age, $D(x, t)$, not just the yearly total $D(t)$.

```
get_ine_deaths(table_id = 6545, age_table_id = 6547, n_periods = 100,
  use_cache = TRUE, force = FALSE,
  cache_dir = tools::R_user_dir("inedemogR", "cache"))
```

Returns `list(data_provinces, data_national, qc, qc_age)`: `data_provinces` is the age-specific series (`nuts3_code`, `province_name`, `year`, `age`, `female`, `male`, `total`) that feeds the rest of the mortality pipeline; `data_national` is Spain’s age-less annual total, used internally to cross-check the sum of provincial totals against the official national figure (`validate_deaths()`’s `national_vs_provincial_mismatch` check) and also validates the sex ratio at death lies in the plausible range $[0.90, 1.50]$.

```
deaths <- get_ine_deaths()
deaths$data_provinces # age-specific, feeds compute_exposure()/compute_death_rates()
deaths$data_national  # Spain total, used for QC cross-checks
```

2.3 The mortality pipeline

The three functions below implement, in order, Steps 4–6 of the SHMD Protocol (equivalently, HMD Methods Protocol V6): exposure-to-risk, central death rates, and period life tables. Each stage consumes the previous stage’s output directly.

2.3.1 Exposure-to-risk: compute_exposure()

Exposure-to-risk $E(x, t)$ approximates the person-years lived at age x during calendar year t , combining the population stock at the start and end of the year with a Lexis-triangle correction for within-year deaths:

$$E(x, t) = \frac{1}{2} [P(x, t) + P(x, t + 1)] + \frac{1}{2} [D^U(x, t) - D^L(x, t)] \quad (1)$$

where $P(x, t)$ is the January 1 population aged x in year t , and $D^U(x, t)$, $D^L(x, t)$ are deaths in the upper and lower Lexis triangles of the (x, t) parallelogram — i.e. deaths to the birth cohort $t - x$ (upper triangle, born before the start of year t , dying at age x) versus deaths to the birth cohort $t - x - 1$ (lower triangle, born during year t itself among those who reach age x). INE’s Tempus3 death tables do not carry a birth-cohort split, so absent a cohort column, `compute_exposure()` falls back to the documented HMD Methods Protocol V6 Appendix A approximation of an even 50/50 split ($D^U = D^L = D(x, t)/2$, which cancels the correction term entirely, i.e. reduces to a plain average of the two population stocks).

```
compute_exposure(population, deaths)
```

population `get_ine_population()`\$data.

```
deaths      get_ine_deaths()$data_provinces.
```

Because death registration lags population estimates by roughly a year, the most recent year in population commonly has no matching rows in deaths yet. That year is still used as the $P(x, t + 1)$ boundary needed to compute the *previous* year's exposure, but is not itself returned as an exposure row — treating an entirely-absent year as zero deaths would mechanically produce zero mortality and, downstream, a nonsensical inflated life expectancy.

```
pop <- get_ine_population()
deaths <- get_ine_deaths()
exposure <- compute_exposure(pop$data, deaths$data_provinces)
exposure$qc$passed
```

`validate_exposure()` flags cells with no usable source population (informational), any remaining negative/NA exposure, and implausible exposure exceeding twice the source population (a generous plausibility bound).

2.3.2 Central death rates: `compute_death_rates()`

Given deaths $D(x, t)$ and exposure $E(x, t)$, the central death rate at single-year age x in year t is

$$m(x, t) = \frac{D(x, t)}{E(x, t)} \quad (1 \times 1: \text{single-year age, one-year period}) \quad (2)$$

INE's age-specific deaths table only publishes rows for age $\times$ year cells with at least one death; `compute_death_rates()` starts from the complete 0–MAX_AGE exposure grid and left-joins deaths, filling absent rows with zero deaths rather than silently dropping them — a missing INE row means *zero* deaths at that age, not a data gap, and dropping it would break the age-contiguity that the life-table recursion (Section 2.3.3) requires.

A pooled 5-year-age-group ($n = 5$) rate is also computed:

$$m(x, n, t) = \frac{\sum_{k=0}^4 D(x + k, t)}{\sum_{k=0}^4 E(x + k, t)} \quad (5 \times 1) \quad (3)$$

and an age-standardized death rate using the Eurostat European Standard Population 2013 (ESP2013) single-year age weights $w(x)$, with $\sum_x w(x) = 100,000$:

$$\text{ASDR}(t) = \sum_x m(x, t) w(x) \quad (4)$$

expressed in deaths per 100,000 of the ESP2013 standard population, allowing direct comparison of mortality levels across provinces or years with different underlying age structures.

```
compute_death_rates(deaths, exposure)
```

Returns `list(mx_1x1, mx_5x1, asdr, qc)`. `validate_death_rates()` flags negative/NA/infinite rates and any rate ≥ 1 below age 90 (implausible for a working-age population, informational).

```
rates <- compute_death_rates(deaths$data_provinces, exposure$data)
rates$mx_1x1 # age x year x province mortality schedule
rates$asdr   # age-standardized rate, one row per province-year
```

2.3.3 Period life tables: `build_life_tables()`

`build_life_tables()` constructs standard 1x1 period life tables for every province and year in `mx_1x1`, following HMD Methods Protocol V6 exactly. The construction proceeds in two stages: old-age smoothing, then the life-table recursion itself.

Kannisto smoothing of old-age mortality. Single-year mortality rates at the oldest ages are noisy (small populations, few deaths), so ages 80–110 are replaced with a smooth Kannisto model fitted on ages 80–95 and extrapolated to age 110:

$$\text{logit } m(x) = \log \frac{m(x)}{1 - m(x)} = \alpha + \beta x, \quad x \in [80, 95] \quad (5)$$

fitted by ordinary least squares, then evaluated for every age up to the open-ended terminal age (110):

$$m^{\text{Kannisto}}(x) = \frac{e^{\alpha + \beta x}}{1 + e^{\alpha + \beta x}} \quad (6)$$

Andreev–Kingkade a_0 . The average person-years lived by those who die before age 1, a_0 , uses the Andreev–Kingkade (2015) piecewise regression on the age-0 central death rate m_0 (replacing the older Coale–Demeny approximation), separately by sex:

$$a_0 = \begin{cases} 0.330 & \text{if } m_0 \geq 0.107 \\ 0.053 + 2.800 m_0 & \text{if female and } m_0 < 0.107 \\ 0.045 + 2.684 m_0 & \text{if male and } m_0 < 0.107 \end{cases} \quad (7)$$

For the combined-sex table, a_0 is the mean of the female and male values. The remaining a_x values follow the standard HMD convention: $a_1 = 0.4$ (age 1), $a_x = 0.5$ for ages 2 and above.

Life-table recursion. Given m_x and a_x for every age $x = 0, 1, \dots, \omega$ (open interval at $\omega = 110$):

$$q_x = \frac{m_x}{1 + (1 - a_x)m_x}, \quad q_\omega = 1 \quad (8)$$

$$l_0 = 100,000 \quad (9)$$

$$d_x = l_x q_x \quad (10)$$

$$l_{x+1} = l_x - d_x \quad (11)$$

$$L_x = \begin{cases} l_{x+1} + a_x d_x & x < \omega \\ l_\omega / m_\omega & x = \omega \end{cases} \quad (12)$$

$$T_x = \sum_{y \geq x} L_y \quad (13)$$

$$e_x = T_x / l_x \quad (14)$$

The open-interval formula $L_\omega = l_\omega / m_\omega$ assumes ω is genuinely old (a high m_ω); this is the reason the terminal age must be a realistic $\omega = 110$ after Kannisto extrapolation rather than an artificially low cut-off, which would otherwise produce a nonsensical remaining-life-expectancy figure in the thousands of years.

```
build_life_table(mx_df, sex = c("female", "male")) # single year, single sex
build_life_tables(mx_1x1) # every province and year
```

`build_life_tables()` returns `list(fltper, mltper, bltper, qc)` — female, male, and both-sex period life tables, each a tibble with `nuts3_code`, `province_name`, `year`, `age`, `mx`, `qx`, `ax`, `lx`, `dx`, `Lx`, `Tx`, `ex`. `validate_life_table()` checks that l_x is non-increasing in age, $L_x > 0$, $e_x > 0$, and $q_x \in [0, 1]$.

```
lt <- build_life_tables(rates$mx_1x1)
lt$fltper[lt$fltper$age == 0, ] # e0 by province and year, female
```

As a live-verified example, this pipeline gives a period life expectancy at birth for A Coruña, 2024, of $e_0 = 86.97$ years (female) and 81.21 years (male) — plausible, correct figures for a Spanish province in that year.

2.3.4 Abridged life tables: `build_abridged_life_tables()`

`build_abridged_life_tables()` builds standard abridged (5-year age group) period life tables directly from `compute_death_rates()`'s `mx_5x1`, complementing `build_life_tables()`'s single-year (1x1) tables — useful for comparing against other agencies' published abridged tables, or when only grouped-age data is available. Age groups are "00-04", "05-09", ..., "95-99", "100+" (open interval).

The one exception to building strictly from `mx_5x1` is the youngest group: INE's own age grouping merges age 0 with ages 1–4 into a single "00-04" group, which would otherwise need an external, less-precise ${}_5a_0$ regression (e.g. Coale–Demeny) to handle the sharp within-interval infant-mortality gradient. Since `mx_1x1` already carries the exact single-year m_0 with its Andreev–Kingkade a_0 (Section 2.3.3), `build_abridged_life_table()` instead builds an exact single-year sub-table for ages 0–4 and aggregates it, avoiding that approximation for the one interval where it matters most.

For every other closed interval $[x, x+n)$ ($n = 5$), the standard discrete abridged life-table identities apply (Preston, Heuveline & Guillot, 2001, Ch. 3):

$${}_na_x = n/2 \quad (15)$$

$${}_nq_x = 1 - e^{-{}_nm_x} \quad (16)$$

The exponential form for ${}_nq_x$ (rather than the linear Chiang formula ${}_nm_x/(1 + (n - {}_na_x){}_nm_x)$ that `build_life_table()` uses for single-year, $n = 1$ intervals) is deliberate: the linear formula can exceed 1 once n is as wide as 5 and mortality is high (old-age groups with ${}_nm_x \gtrsim 0.3$ –0.4), which the exponential form never does. The remaining recursion matches Section 2.3.3 exactly, with the person-years-lived identity generalized to the interval width:

$${}_nL_x = n l_{x+n} + {}_na_x n d_x \quad (x < \omega) \quad (17)$$

```
build_abridged_life_table(mx_1x1, mx_5x1, sex = c("female", "male")) # one province/year
build_abridged_life_tables(mx_1x1, mx_5x1) # every province/year
```

Returns `list(fltper, mltper, bltper, qc)`, matching `build_life_tables()`'s structure with `age_group/age_start` in place of `age`. `validate_abridged_life_table()` checks the same invariants as `validate_life_table()`.

```
alt <- build_abridged_life_tables(rates$mx_1x1, rates$mx_5x1)
alt$fltper[alt$fltper$age_group == "00-04", ]
```

As a live-verified cross-check, A Coruña 2024's abridged e_0 (female) comes to 87.03 years against the exact single-year figure of 86.97 — a 0.06-year difference, consistent with the coarser 5-year grouping.

2.4 Exporting to files: `download_ine_data()`

For users who prefer to work outside R (a spreadsheet, another statistical package, or HMD-compatible mortality software), `download_ine_data()` runs any subset of the six pipeline stages and writes the results to a folder as CSV and/or HMD-format `.txt` files.

```
download_ine_data(out_dir,
  stages = c("births", "deaths", "population",
             "exposure", "mx", "life_tables"),
  format = c("csv", "txt"),
  n_periods = 30)
```

Later stages automatically pull in their dependencies (exposure needs population and deaths; mx needs deaths and exposure; life_tables needs mx), since they are fetched along the way regardless.

```
# Everything, both file formats, into ./ine_data
result <- download_ine_data("ine_data")
result$files          # every path written
result$life_tables$fltper # also available in memory, no re-read needed

# Just births and deaths, CSV only
download_ine_data("ine_data", stages = c("births", "deaths"), format = "csv")

# One HMD-format .txt file per province
download_ine_data("ine_data", stages = "life_tables", format = "txt")
```

"csv" writes one combined file per data table (e.g. `population.csv`, one row per province×age×year); "txt" writes HMD-format files, one per province, matching the layout HMD-compatible tooling expects.

3 Part II — Demographic Analysis and Visualization

3.1 Summary demographic indicators

The functions in this section all operate on System B's age/sex-disaggregated province data (or, for `birth_death_ratio()`, on System A's totals — noted explicitly below); `get_ine_demog()`'s total-only indicators cannot be used for any indicator that requires an age breakdown.

3.1.1 Age dependency ratios: `age_dependency_ratio()`

With P_{young} , P_{work} , P_{old} the population aged ≤ 14 , 15–64, and ≥ 65 respectively (cut-offs configurable via `young_max/old_min`):

$$\text{youth dependency ratio} = \frac{P_{\text{young}}}{P_{\text{work}}} \times 100 \quad (18)$$

$$\text{old-age dependency ratio} = \frac{P_{\text{old}}}{P_{\text{work}}} \times 100 \quad (19)$$

$$\text{total dependency ratio} = \frac{P_{\text{young}} + P_{\text{old}}}{P_{\text{work}}} \times 100 \quad (20)$$

```
pop <- get_ine_population()
adr <- age_dependency_ratio(pop$data)
adr[adr$nuts3_code == "ES111", ]
```

3.1.2 Aging index: `aging_index()`

$$\text{aging index} = \frac{P_{\text{old}}}{P_{\text{young}}} \times 100 \quad (21)$$

```
ai <- aging_index(pop$data)
```

3.1.3 Sex ratio: `sex_ratio()`

$$\text{sex ratio} = \frac{P_{\text{male}}}{P_{\text{female}}} \times 100 \quad (22)$$

computed over all ages by default, or per single year of age with `by_age = TRUE`.

```
sex_ratio(pop$data)
sex_ratio(pop$data, by_age = TRUE)
```

3.1.4 Crude birth and death rates

$$\text{CBR} = \frac{B}{P} \times 1000, \quad \text{CDR} = \frac{D}{P} \times 1000 \quad (23)$$

where B , D , P are total annual births, deaths, and population.

```
births <- get_ine_births()
deaths <- get_ine_deaths()
cbr <- crude_birth_rate(births$data, pop$data)
cdr <- crude_death_rate(deaths$data_provinces, pop$data)
```

These are crude rates, not a total fertility rate (TFR). TFR requires age-of-mother-specific birth counts; INE’s underlying MNPN table as retrieved by `get_ine_births()` carries no age-of-mother breakdown, and this package does not currently compute TFR. Do not substitute `crude_birth_rate()` for TFR in fertility analysis.

3.1.5 General fertility rate: `general_fertility_rate()`

GFR improves on CBR by dividing births by the female population actually at risk of childbearing (ages 15–49 by default), removing the distortion CBR suffers when two provinces have different age/sex structures despite similar underlying fertility behavior:

$$\text{GFR} = \frac{B}{W_{15-49}} \times 1000 \quad (24)$$

```
gfr <- general_fertility_rate(births$data, pop$data, age_min = 15, age_max = 49)
```

GFR is still not a TFR: it is a single aggregate rate over all reproductive-age women, not an age-specific fertility schedule.

3.1.6 Fertility schedule: ASFR, TFR, MAC, GRR, NRR

Where CBR and GFR (above) are single aggregate rates, the functions below use `get_ine_births_by_age()`’s (Section 2.2.4) age-of-mother breakdown to build the full age-specific fertility schedule and the standard summary indicators derived from it.

Age-specific fertility rate: `age_specific_fertility_rate()`.

$$\text{ASFR}(x) = \frac{B(x)}{W(x)} \times 1000, \quad \text{ASFR}_{\text{female}}(x) = \frac{B_{\text{female}}(x)}{W(x)} \times 1000 \quad (25)$$

for single years of age $x = 15, \dots, 49$ by default, where $B(x)$ is births to mothers of age x and $W(x)$ is the female population of age x . `age_specific_fertility_rate()` returns both the all-newborn rate $\text{ASFR}(x)$ and the female-newborn-specific rate $\text{ASFR}_{\text{female}}(x)$ (using INE’s actual sex-of-newborn breakdown, not an assumed constant sex ratio at birth) — the schedule the remaining functions below all consume.

```
births_age <- get_ine_births_by_age()
pop <- get_ine_population()
asfr <- age_specific_fertility_rate(births_age$data, pop$data, age_min = 15, age_max = 49)
```

3.1.7 Total fertility rate: `total_fertility_rate()`

$$\text{TFR} = \sum_{x=15}^{49} \frac{\text{ASFR}(x)}{1000} \quad (26)$$

the true TFR that `crude_birth_rate()` and `general_fertility_rate()` explicitly note this package could not previously compute.

```
tfr <- total_fertility_rate(asfr)
```

3.1.8 Mean age at childbearing: `mean_age_at_childbearing()`

$$\text{MAC} = \frac{\sum_x x \cdot \text{ASFR}(x)}{\sum_x \text{ASFR}(x)} \quad (27)$$

```
mac <- mean_age_at_childbearing(asfr)
```

3.1.9 Gross and net reproduction rates

The gross reproduction rate (expected daughters per woman over her reproductive lifetime, ignoring mortality) sums the female-newborn-specific schedule directly, rather than approximating it by scaling TFR with an assumed constant sex ratio at birth:

$$\text{GRR} = \sum_{x=15}^{49} \frac{\text{ASFR}_{\text{female}}(x)}{1000} \quad (28)$$

The net reproduction rate additionally weights that schedule by the mother's own survivorship to each reproductive age, using the female life table's L_x (Section 2.3.3) and radix l_0 :

$$\text{NRR} = \sum_{x=15}^{49} \frac{\text{ASFR}_{\text{female}}(x)}{1000} \cdot \frac{L_x}{l_0} \quad (29)$$

so $\text{NRR} \leq \text{GRR}$ always, with equality only in the (unrealistic) absence of any mortality before the end of the reproductive span.

```
grr <- gross_reproduction_rate(asfr)
nrr <- net_reproduction_rate(asfr, lt$fltptr) # lt$fltptr: female life table, build_life_tables()
```

As a live-verified example, Madrid 2023 gives $\text{GRR} = 0.538$ and $\text{NRR} = 0.535$ — both well below the replacement threshold of 1, consistent with Spain's well-documented below-replacement fertility, with NRR correctly just under GRR given Spain's low mortality at reproductive ages.

3.1.10 Infant mortality rate: `infant_mortality_rate()`

$$\text{IMR} = \frac{D_0}{B} \times 1000 \quad (30)$$

where D_0 is deaths under age 1 in year t and B is live births in that *same* year t (the birth cohort those deaths are drawn from) — a different denominator convention from CDR. IMR is closely related to, but not identical to, q_0 in the period life table: q_0 is derived from the exposure-based m_0 via the Andreev-Kingkade a_0 , while IMR here uses the simpler, standard deaths/births ratio directly; the two will usually be close but need not match exactly.

```
imr <- infant_mortality_rate(deaths$data_provinces, births$data)
```

3.1.11 Rate of natural increase: `rate_of_natural_increase()`

$$\text{RNI} = \text{CBR} - \text{CDR} \quad (31)$$

combining already-computed `crude_birth_rate()` and `crude_death_rate()` output. A positive value means the population is growing from natural change alone (ignoring migration); negative means the reverse — the point in time and space where CDR overtakes CBR is often called the “demographic crossover.”

```
rni <- rate_of_natural_increase(cbr, cdr)
```

3.1.12 Birth-to-death ratio: `birth_death_ratio()`

Unlike the indicators above, this needs no age breakdown, so it operates directly on `get_ine_demog()`’s totals (System A):

$$\text{birth-death ratio} = \frac{B}{D} \quad (32)$$

```
vitals <- get_ine_demog(indicator = c("births_total", "deaths_total"), year = 2023)
ratio <- birth_death_ratio(vitals)
```

A value of 2 means 2 live births per death; 0.5 means 1 birth per 2 deaths.

3.1.13 Life expectancy: `life_expectancy_summary()` and `life_expectancy()`

`life_expectancy_summary()` extracts e_0 (life expectancy at birth) and e_{65} (remaining life expectancy at 65) for every province and year from a `build_life_tables()` output table:

```
le <- life_expectancy_summary(lt$fltptr, sex = "female")
```

`life_expectancy()` is a one-province convenience wrapper: it builds the life table for a single matched province directly from `compute_death_rates()`’s `mx_1x1` output, without requiring `build_life_tables()` to be run for every province first.

```
life_expectancy(rates$mx_1x1, province = "A Coruna", sex = "female")
```

3.2 Mortality decomposition: `decompose_life_expectancy()`

Life expectancy summaries answer *what* e_0 is; mortality decomposition answers *why* two life tables’ e_0 values differ — attributing the gap between two provinces (same year) or one province (two years) to age-specific contributions. Two independently-derived methods are implemented so results can be cross-checked against each other (Ponnapalli, 2005, found the two are not sensitive to which one is used):

```
decompose_life_expectancy(lt1, lt2, method = c("arriaga", "pollard"))
```

lt1, lt2 Single life tables (one province-year-sex slice of `build_life_tables()`'s `fltper/mltper/bltper`, or `build_life_table()`'s direct output), over the *same* set of ages.

Arriaga's method (1984, the default). An exact discrete decomposition. For a non-terminal age interval $[x, x+n)$, splitting into a direct effect (extra/fewer years lived *within* the interval) and a combined indirect/interaction effect (extra/fewer survivors carrying forward into *later* ages):

$$DE(x) = \frac{l_x^{(1)}}{l_0^{(1)}} \left(\frac{nL_x^{(2)}}{l_x^{(2)}} - \frac{nL_x^{(1)}}{l_x^{(1)}} \right) \quad (33)$$

$$IE(x) = \frac{T_{x+n}^{(2)}}{l_0^{(1)}} \left(\frac{l_x^{(1)}}{l_x^{(2)}} - \frac{l_{x+n}^{(1)}}{l_{x+n}^{(2)}} \right) \quad (34)$$

$$\Delta(x) = DE(x) + IE(x) \quad (35)$$

with superscripts (1)/(2) denoting the reference/comparison life table. For the terminal open interval ω , only a direct effect applies (since ${}_{\infty}L_{\omega} = T_{\omega}$):

$$\Delta(\omega) = \frac{l_{\omega}^{(1)}}{l_0^{(1)}} \left(\frac{T_{\omega}^{(2)}}{l_{\omega}^{(2)}} - \frac{T_{\omega}^{(1)}}{l_{\omega}^{(1)}} \right) \quad (36)$$

and $\sum_x \Delta(x) = e_0^{(2)} - e_0^{(1)}$ exactly.

Pollard's method (1988). An exact *continuous*-time decomposition; applied to a discrete single-year life table, its symmetric-average form (averaging the two individually-exact continuous identities Pollard derives, so the choice of reference table doesn't bias the result) is:

$$\Delta(x) = \frac{1}{2} (m_x^{(1)} - m_x^{(2)}) \left(\frac{l_x^{(2)}}{l_0^{(2)}} e_x^{(1)} + \frac{l_x^{(1)}}{l_0^{(1)}} e_x^{(2)} \right) \quad (37)$$

Unlike Arriaga's, this is exact only in the continuous limit — on a single-year life table it is a close but not identical approximation of $e_0^{(2)} - e_0^{(1)}$, so treat Arriaga's (the default) as the primary result and Pollard's as a cross-check.

```
lt <- build_life_tables(rates$mx_1x1)
lt_a <- lt$fltper[lt$fltper$province_name == "A Coruña" & lt$fltper$year == 2023, ]
lt_b <- lt$fltper[lt$fltper$province_name == "Madrid" & lt$fltper$year == 2023, ]

decomp <- decompose_life_expectancy(lt_a, lt_b, method = "arriaga")
sum(decomp$contribution) # recovers the e0 gap exactly
```

As a live-verified example, A Coruña vs. Madrid, 2023: e_0 gap = 1.04 years, with age 0 and age 58 the largest single-age contributors (Arriaga and Pollard correlate at 0.999 on this comparison). For a temporal comparison — A Coruña, 2016 vs. 2023 — the gap is 0.69 years.

3.3 Mapping functions

3.3.1 plot_ine_map()

Plots the full map of Spain at province or municipality level, with one or more regions highlighted — primarily a diagnostic tool for confirming which regions a regex region/highlight pattern actually matches.

```
plot_ine_map(geo_level = "province", highlight = "^Cordoba$")
```

3.3.2 map_indicator()

General-purpose choropleth: joins any tidy indicator tibble keyed by GEOID onto province or municipality geometry, as either a binned discrete-legend map (default, matching the classic “N regions above/below a threshold” choropleth style) or a continuous gradient.

```
map_indicator(df, value_col, geo_level = c("province", "municipality"),
  region = NULL, binned = TRUE, breaks = NULL,
  palette = if (binned) "PiYG" else "D",
  moveCAN = TRUE, can_gap_km = 60,
  legend_title = NULL, title = NULL)
```

```
pop_data <- get_ine_demog(indicator = "population_total", year = 2023)
map_indicator(pop_data, "population_total", geo_level = "province")
```

3.3.3 map_life_expectancy()

A thin wrapper around map_indicator() that bridges System B’s life_expectancy_summary() output with System A’s get_ine_geo() geometry via the internal province lookup table.

```
map_life_expectancy(le_df, year, sex = "both",
  moveCAN = TRUE, can_gap_km = 60, title = NULL)
```

```
map_life_expectancy(le, year = max(le$year), sex = "female")
```

3.4 Population pyramids and trend charts

3.4.1 plot_population_pyramid()

Draws a mirrored population pyramid (female right, male left as negative counts) for one province/region and year.

```
plot_population_pyramid(pop_df, year, region = NULL, title = NULL)
```

```
plot_population_pyramid(pop$data, year = max(pop$data$year), region = "A Coruna")
```

Note geom_col(width = 1) is used internally rather than ggplot2’s default width = 0.9: at the default width, the systematic 10% gap between each of the 101 single-year-age bars renders as visible thin white stripes across the pyramid.

3.4.2 plot_demog_trend()

A generic time-series chart for any indicator tibble keyed by nuts3_code/province_name/year — works unmodified on the output of any indicator function above.

```
plot_demog_trend(df, value_col, region = NULL, title = NULL, ylab = NULL)
```

```
plot_demog_trend(ai, "aging_index", region = "A Coruna")
```

3.5 Lexis diagrams: plot_lexis_diagram()

A Lexis diagram plots age (vertical axis) against calendar year (horizontal axis), shaded by mortality rate, with birth-cohort diagonals overlaid. Every point along one diagonal line represents the same birth cohort, aging one year for every calendar year that passes:

$$\text{age} = \text{year} - \text{cohort} \quad (38)$$

```
plot_lexis_diagram(mx_df, province, sex = c("total", "female", "male"),  
  log_scale = TRUE, rate_per = 1000,  
  cohort_lines = TRUE, cohort_step = 10, title = NULL)
```

```
rates <- compute_death_rates(deaths$data_provinces, exposure$data)  
plot_lexis_diagram(rates$mx_1x1, province = "A Coruna")  
plot_lexis_diagram(rates$mx_1x1, province = "Madrid", cohort_lines = FALSE)
```

Mortality is color-scaled on a log10 axis by default (`log_scale = TRUE`), the standard convention for Lexis surfaces since m_x spans several orders of magnitude across ages. Grey cells mean no rate could be computed for that age×year cell — almost always because the source population data has no single-year age breakdown for the oldest ages in the earliest years of the series (so exposure, the rate’s denominator, is missing there), not a defect in the calculation; this typically appears as a rectangular block in the upper-left (oldest ages, earliest years) and clears once the source data’s age detail becomes complete. When present, grey is given its own “Missing data” legend key via a small dummy point layer, since a continuous color scale cannot show an `na.value` swatch on its own.

3.6 Worked examples

The nine scripts summarized below live in `examples/` and are fully runnable end to end (each retrieves live data with the package’s own functions before plotting). They double as integration tests for the indicator and charting functions documented above.

3.6.1 Birth-to-death ratio map (birth_death_ratio_map.R)

Replicates a reference choropleth style (binned, pink-to-green diverging scale) at the finest granularity Spain’s public birth/death data actually supports: province (municipality-level birth/death tables do not exist in INE’s MNP/MNPD operations).

```
library(inedemogR)  
  
vitals <- get_ine_demog(indicator = c("births_total", "deaths_total"), year = 2023)
```

```
ratio <- birth_death_ratio(vitals)

map_indicator(
  ratio, value_col = "birth_death_ratio", geo_level = "province",
  breaks = c(0.25, 0.5, 0.75, 1, 1.5, 2, 3), palette = "PiYG",
  legend_title = "Births per death",
  title = "Birth-to-death ratio by province, Spain (2023)"
)
```

3.6.2 Andalusian population pyramids (andalusia_pyramids.R)

Builds an eight-panel `facet_wrap()` grid of population pyramids for Andalusia's eight provinces (all sharing `nuts2_code == "ES61"`), directly in **ggplot2** syntax rather than calling `plot_population_pyramid()` eight times.

```
library(inedemogR); library(dplyr); library(tidyr); library(ggplot2)

pop <- get_ine_population()
latest_year <- max(pop$data$year)
andalusia <- pop$data |> filter(nuts2_code == "ES61", year == latest_year)

pyramid_data <- andalusia |>
  select(province_name, age, female, male) |>
  pivot_longer(c(female, male), names_to = "sex", values_to = "count") |>
  mutate(count = if_else(sex == "male", -count, count))

ggplot(pyramid_data, aes(x = age, y = count, fill = sex)) +
  geom_col(width = 1) +
  coord_flip() +
  facet_wrap(~province_name, nrow = 2, ncol = 4, scales = "free_x") +
  scale_y_continuous(labels = abs) +
  scale_fill_manual(values = c(female = "#D4667A", male = "#4477AA")) +
  labs(x = "Age", y = "Population", fill = "Sex",
       title = paste("Population pyramids, Andalusian provinces,", latest_year)) +
  theme_minimal()
```

3.6.3 Madrid/Barcelona Lexis diagrams (madrid_barcelona_lexis.R)

`plot_lexis_diagram()` handles one province per call; a two-panel comparison is built directly with the pipeline functions plus `facet_wrap()`.

```
library(inedemogR); library(dplyr); library(ggplot2)

pop <- get_ine_population()
deaths <- get_ine_deaths()
exposure <- compute_exposure(pop$data, deaths$data_provinces)
rates <- compute_death_rates(deaths$data_provinces, exposure$data)

mx <- rates$mx_1x1 |> filter(province_name %in% c("Madrid", "Barcelona")) |>
  mutate(mx_per_1000 = mx_total * 1000)

ggplot(mx, aes(x = year, y = age, fill = mx_per_1000)) +
  geom_tile() +
  scale_fill_viridis_c(
    name = "Mortality rate\n(per 1,000, log scale)", trans = "log10",
```

```

na.value = "grey50", labels = scales::label_number(accuracy = 0.1)
) +
facet_wrap(~province_name, nrow = 1, ncol = 2) +
labs(x = "Year", y = "Age", title = "Lexis diagrams: Madrid and Barcelona") +
theme_minimal()

```

3.6.4 Life expectancy comparison (life_expectancy_comparison.R)

Combines a multi-province time series with a national choropleth for one year, both from the same life-table pipeline run once for every province.

```

library(inedemogR); library(dplyr); library(ggplot2)

pop <- get_ine_population()
deaths <- get_ine_deaths()
exposure <- compute_exposure(pop$data, deaths$data_provinces)
rates <- compute_death_rates(deaths$data_provinces, exposure$data)
lt <- build_life_tables(rates$mx_1x1)
le <- life_expectancy_summary(lt$fltper, sex = "female")

provinces_to_compare <- c("Madrid", "Barcelona", "Sevilla", "A Coruna")
ggplot(le |> filter(province_name %in% provinces_to_compare),
       aes(x = year, y = e0, color = province_name)) +
  geom_line(linewidth = 0.8) + geom_point(size = 1.5) +
  labs(x = "Year", y = "Life expectancy at birth (e0)", color = "Province") +
  theme_minimal()

map_life_expectancy(le, year = max(le$year), sex = "female",
                   title = "Life expectancy at birth by province, Spain")

```

3.6.5 Natural increase analysis (natural_increase_analysis.R)

Visualizes the “demographic crossover” (CDR overtaking CBR) and the resulting rate of natural increase.

```

library(inedemogR); library(dplyr); library(ggplot2)

births <- get_ine_births(); deaths <- get_ine_deaths(); pop <- get_ine_population()
cbr <- crude_birth_rate(births$data, pop$data)
cdr <- crude_death_rate(deaths$data_provinces, pop$data)
rni <- rate_of_natural_increase(cbr, cdr)

provinces_to_compare <- c("Madrid", "Barcelona", "Sevilla", "A Coruna")
ggplot(rni |> filter(province_name %in% provinces_to_compare),
       aes(x = year, y = rni, color = province_name)) +
  geom_hline(yintercept = 0, linetype = "dashed", color = "grey50") +
  geom_line(linewidth = 0.8) + geom_point(size = 1.5) +
  labs(x = "Year", y = "Rate of natural increase (per 1,000)", color = "Province") +
  theme_minimal()

```

3.6.6 Fertility and infant mortality analysis (fertility_infant_mortality_analysis.R)

```

library(inedemogR); library(dplyr); library(ggplot2)

births <- get_ine_births(); deaths <- get_ine_deaths(); pop <- get_ine_population()

```

```

gfr <- general_fertility_rate(births$data, pop$data)
imr <- infant_mortality_rate(deaths$data_provinces, births$data)

provinces_to_compare <- c("Madrid", "Barcelona", "Sevilla", "A Coruna")
ggplot(gfr |> filter(province_name %in% provinces_to_compare),
  aes(x = year, y = gfr, color = province_name)) +
  geom_line(linewidth = 0.8) + geom_point(size = 1.5) +
  labs(x = "Year", y = "General fertility rate (per 1,000 women 15-49)",
    color = "Province") +
  theme_minimal()

```

3.6.7 Fertility indicators analysis (fertility_indicators_analysis.R)

Exercises the full ASFR-based fertility schedule (Section 3.1.6): the age schedule itself, TFR/MAC trends, and the GRR-vs-NRR mortality-adjustment gap, for a comparison set of provinces.

```

library(inedemogR); library(dplyr); library(ggplot2)

births_age <- get_ine_births_by_age()
pop <- get_ine_population()
deaths <- get_ine_deaths()
asfr <- age_specific_fertility_rate(births_age$data, pop$data)
tfr <- total_fertility_rate(asfr)
mac <- mean_age_at_childbearing(asfr)
grr <- gross_reproduction_rate(asfr)

exposure <- compute_exposure(pop$data, deaths$data_provinces)
rates <- compute_death_rates(deaths$data_provinces, exposure$data)
lt <- build_life_tables(rates$mx_1x1)
nrr <- net_reproduction_rate(asfr, lt$fltper)

latest_year <- max(asfr$year)
provinces_to_compare <- c("Madrid", "Barcelona", "Sevilla", "A Coruna")

ggplot(asfr |> filter(province_name %in% provinces_to_compare, year == latest_year),
  aes(x = age, y = asfr, color = province_name)) +
  geom_line(linewidth = 0.8) + geom_point(size = 1.2) +
  labs(x = "Age of mother", y = "ASFR (births per 1,000 women of that age)",
    color = "Province", title = paste("ASFR schedule,", latest_year)) +
  theme_minimal()

ggplot(tfr |> filter(province_name %in% provinces_to_compare),
  aes(x = year, y = tfr, color = province_name)) +
  geom_line(linewidth = 0.8) + geom_point(size = 1.5) +
  geom_hline(yintercept = 2.1, linetype = "dashed", color = "grey40") +
  labs(x = "Year", y = "Total fertility rate", color = "Province",
    subtitle = "Dashed line: replacement-level TFR (2.1)") +
  theme_minimal()

```

Live-verified results: TFR ranges roughly 1.0–1.3 across the four provinces (Madrid highest, A Coruña lowest), well below replacement, and Madrid’s mean age at childbearing comes to 32.8 years — consistent with Spain’s known late-childbearing pattern.

3.6.8 Mortality decomposition analysis (mortality_decomposition_analysis.R)

Exercises `decompose_life_expectancy()` (Section 3.2) on both a cross-sectional gap (two provinces, one year) and a temporal gap (one province, two years), comparing Arriaga's and Pollard's methods.

```
library(inedemogR); library(dplyr); library(ggplot2)

pop <- get_ine_population(); deaths <- get_ine_deaths()
exposure <- compute_exposure(pop$data, deaths$data_provinces)
rates <- compute_death_rates(deaths$data_provinces, exposure$data)
lt <- build_life_tables(rates$mx_1x1)

latest_year <- max(lt$fltper$year); earliest_year <- min(lt$fltper$year)
lt_a <- lt$fltper |> filter(province_name == "A Coruna", year == latest_year)
lt_b <- lt$fltper |> filter(province_name == "Madrid", year == latest_year)
decomp <- decompose_life_expectancy(lt_a, lt_b, method = "arriaga")

ggplot(decomp, aes(x = age, y = contribution, fill = contribution > 0)) +
  geom_col(width = 1) +
  scale_fill_manual(values = c(TRUE = "steelblue", FALSE = "firebrick"), guide = "none") +
  labs(x = "Age", y = "Contribution to e0 gap (years)",
       title = "Madrid vs. A Coruna life expectancy gap by age") +
  theme_minimal()
```

Live-verified: A Coruña 2024 $e_0 = 86.97$ (matching the value confirmed in Section 2.3.3 exactly), Madrid 2024 $e_0 = 87.83$ — gap 0.86 years; the temporal comparison (A Coruña, 1996–2024) gives a gap of 5.23 years, consistent with the secular mortality improvement documented throughout this tutorial.

3.6.9 Abridged life tables analysis (abridged_life_tables_analysis.R)

Cross-checks `build_abridged_life_tables()` (Section 2.3.4) against the exact single-year `build_life_tables()` on the same underlying rates, then charts the abridged survivorship curve and mortality schedule.

```
library(inedemogR); library(dplyr); library(ggplot2)

pop <- get_ine_population(); deaths <- get_ine_deaths()
exposure <- compute_exposure(pop$data, deaths$data_provinces)
rates <- compute_death_rates(deaths$data_provinces, exposure$data)

lt_1x1 <- build_life_tables(rates$mx_1x1)
lt_abridged <- build_abridged_life_tables(rates$mx_1x1, rates$mx_5x1)

latest_year <- max(lt_abridged$fltper$year)
lt_1x1_one <- lt_1x1$fltper |> filter(province_name == "A Coruna", year == latest_year)
lt_abridged_one <- lt_abridged$fltper |> filter(province_name == "A Coruna", year == latest_year)

survivorship <- bind_rows(
  lt_1x1_one |> transmute(age = age, lx = lx, table = "Single-year (1x1)"),
  lt_abridged_one |> transmute(age = age_start, lx = lx, table = "Abridged (5x1)")
)

ggplot(survivorship, aes(x = age, y = lx, color = table)) +
  geom_line(linewidth = 0.8) +
  labs(x = "Age", y = "Survivors (out of 100,000 births)", color = NULL) +
  theme_minimal()
```

Live-verified: A Coruña 2024 abridged $e_0 = 87.03$ against the exact single-year 86.97 — a 0.06-year differ-

ence, confirming the abridged construction recovers the single-year result closely.

Appendix

A.1 Function reference

Function	Purpose
<code>get_ine_demog()</code>	System A: totals for one or more indicators, any geography
<code>get_ine_geo()</code>	Province/municipality boundary geometries
<code>list_ine_indicators()</code>	The indicator registry
<code>update_ine_data()</code>	Refresh cached System A indicators if INE published newer data
<code>plot_ine_map()</code>	Highlight-region diagnostic map
<code>map_indicator()</code>	General choropleth for any GEOID-keyed tibble
<code>get_ine_population()</code>	Province population by single-year age/sex
<code>get_ine_births()</code>	Province births by sex
<code>get_ine_births_by_age()</code>	Province births by single year of age of mother and sex of newborn
<code>get_ine_deaths()</code>	Province deaths by sex, age-specific and age-less
<code>compute_exposure()</code>	Exposure-to-risk $E(x, t)$
<code>compute_death_rates()</code>	Central death rates: 1x1, 5x1, ASDR
<code>andreev_kingkade_a0()</code>	Age-0 life-table parameter a_0
<code>build_life_table()</code>	/ Period life table, one sex/one province, or all
<code>build_life_tables()</code>	
<code>build_abridged_life_table()</code>	Abridged (5-year age group) period life table, one sex/one
<code>/build_abridged_life_tables()</code>	province, or all
<code>download_ine_data()</code>	Export the whole pipeline to CSV/HMD-txt files
<code>validate_population(),</code>	Quality-control checks for each pipeline stage
<code>validate_births(),</code>	
<code>validate_births_by_age(),</code>	
<code>validate_deaths(),</code>	
<code>validate_deaths_age(),</code>	
<code>validate_exposure(),</code>	
<code>validate_death_rates(),</code>	
<code>validate_life_table(),</code>	
<code>validate_abridged_life_table()</code>	
<code>age_dependency_ratio()</code>	Youth/old-age/total dependency ratios
<code>aging_index()</code>	Elderly per 100 children
<code>sex_ratio()</code>	Males per 100 females
<code>crude_birth_rate()</code>	/ CBR / CDR
<code>crude_death_rate()</code>	
<code>general_fertility_rate()</code>	GFR (births per 1,000 women 15-49)
<code>age_specific_fertility_rate()</code>	ASFR/ASFR-female schedule by single year of age
<code>total_fertility_rate()</code>	TFR (children per woman)
<code>mean_age_at_childbearing()</code>	MAC
<code>gross_reproduction_rate()</code>	GRR (daughters per woman, ignoring mortality)
<code>net_reproduction_rate()</code>	NRR (daughters per woman, mortality-adjusted)
<code>infant_mortality_rate()</code>	IMR (infant deaths per 1,000 live births)
<code>rate_of_natural_increase()</code>	RNI = CBR – CDR

Function	Purpose
<code>birth_death_ratio()</code>	Births per death, System A totals
<code>life_expectancy_summary()</code>	e_0/e_{65} from a full life-table output
<code>life_expectancy()</code>	e_0/e_{65} , one-province convenience wrapper
<code>decompose_life_expectancy()</code>	Age-specific decomposition of an e_0 gap (Arriaga/Pollard)
<code>plot_population_pyramid()</code>	Mirrored age/sex pyramid
<code>plot_demog_trend()</code>	Generic indicator time series
<code>plot_lexis_diagram()</code>	Age $\times$ year mortality surface with cohort diagonals
<code>map_life_expectancy()</code>	Life-expectancy choropleth (bridges System B to System A geometry)

A.2 INE data sources used

Series	Table ID	Operation
Population (municipality)	29005	Cifras oficiales del padrón (DPOP)
Population (province, age/sex)	56945	Padrón Municipal Continuo (ECP)
Births (province)	6506	Movimiento Natural de la Población — Nacimientos (MNP)
Births (province, age of mother)	6508	Movimiento Natural de la Población — Nacimientos (MNP)
Deaths (province, age-less)	6545	Movimiento Natural de la Población — Defunciones (MNP)
Deaths (province, age-specific)	6547	Movimiento Natural de la Población — Defunciones (MNP)

A.3 Methodology references

- Human Mortality Database. *Methods Protocol for the Human Mortality Database*, Version 6. University of California, Berkeley, and Max Planck Institute for Demographic Research.
- Andreev, E. M., & Kingkade, W. W. (2015). Average age at death in infancy and infant mortality level: Reconsidering the Coale-Demeny formulas at current levels of low mortality. *Demographic Research*, 33, 363–390.
- Eurostat. *Revision of the European Standard Population* (2013 edition).
- Preston, S. H., Heuveline, P., & Guillot, M. (2001). *Demography: Measuring and Modeling Population Processes*. Blackwell Publishing. (Abridged life-table construction, Ch. 3.)
- Arriaga, E. E. (1984). Measuring and explaining the change in life expectancies. *Demography*, 21(1), 83–96.
- Pollard, J. H. (1988). On the decomposition of changes in expectation of life and differentials in life expectancy. *Demography*, 25(2), 265–276.
- Instituto Nacional de Estadística (INE). *Fenómenos demográficos*. <https://www.ine.es>.