



SDC Research at FHNW & swissanon

sdcMicro, simPop, synvey, riskutility — and
AI on both sides of the table

Matthias Templ

School of Business (ICC),
FHNW ■ swissanon — Swiss
Data Anonymization Center

SDC4Society ■ Tokyo, August
2026



University of Applied Sciences and Arts Northwestern Switzerland
School of Business

What I would like to show you

1. **swissanon** — what we are and what we do
2. **The R toolbox** — `sdcMicro`, `simPop`, `synvey`, `riskutility` / RAPID
3. **Currently running projects** — SNF Bridge, Nestlé, SBB (and a long tail of past ones)
4. **AI on both sides of the table** — a new assistant, and a new adversary
5. **Working together** — two concrete ideas

Informal session — please interrupt with questions at any point. The last part is a proposal, not a conclusion.

From an R package to a data-anonymization center

Who


- Matthias Templ, **FHNW** — School of Business (ICC), Olten.
- Statistical disclosure control, synthetic data, robust statistics.
- Building SDC software in R since `sdcMicro` (CRAN, 2007); ~90 peer-reviewed papers.

swissanon — Swiss Data Anonymization Center

- Three pillars: **Research** ■ **Courses** ■ **For Companies**.
- Every method is peer-reviewed, reproducible, and open — “**risk is measured, not asserted**”.
- For companies: disclosure-risk assessment, anonymization, synthetic data, and independent audits.

A center built around SDC research — the same shape as an SDC4Society project. swissanon.com


Two roads from real data



From real data there are two roads — **synthesize** it, or **anonymize** it. Either way the result is **measured before release**: risk is quantified, not asserted.

The next five slides walk through these four packages.

sdcMicro — the workhorse: measure, then treat



Risk → **recode** → **suppress** → **perturb** → utility → **document**. Risk metrics: k -anonymity, individual & population risk, SUDA, ℓ -diversity. Treatments: recoding, local suppression, PRAM, microaggregation, swapping, noise.

CRAN since 2007 (v5.8.2); GUI `sdApp`; the audit report ships *with* the data — *it is* the deliverable.

simPop — synthetic *populations* for SDC

- Simulates a full synthetic **population** from a survey plus known margins (model-based simulation, calibration).
- Used by statistical offices to build population microdata and agent-based inputs.
- Mature (v2.1.4); it anchors the *design-aware* corner of the map below.

Where the synthesizers sit

	size- n (sample)	size- N (population)
design-blind	synthpop	—
design-aware	synvey	simPop

synvey fills the empty design-aware, sample-size cell.

synvey — design-aware synthesis at sample size

- `specify()` declares the **design**: weights, strata, household & person IDs. `target = sample | population | design-preserving weighted sample`.
- `draw = pmm / residual / normal` is a **utility ↔ disclosure dial**; `robsynth()` keeps contaminated training data from propagating.

Why design-awareness matters


Design-*blind* synthesis inherits the full informative-sampling bias (mean bias up to **+1.46** in our study); design-aware synthesis removes it, while regression slopes stay robust either way. Membership-inference $AUC \approx 0.5$ across draws.

Public since Aug 2026 (github.com/matthias-da/synvey); paper under review (Survey Methodology). Suggests riskutility — the next box.


riskutility — one grammar for risk and utility

- ~60 measures under **one S3 API** (`print/summary/plot`), for anonymized *and* synthetic data.
- `rumap()` draws the risk–utility map; `from_simPop()`, `from_syntpop()`, `from_sdcMicro()` adapt any input.
- Families: identity (*k*-anon, SUDA), attribute (CAP/TCAP), **inference (RAPID)**, membership (DOMIAS, NNAA); utility (pMSE, SPECKS, TSTR).
- `disclosure_report()` — one call, a full verdict.

On CRAN since 2026; JSS paper in preparation.



RAPID — risk of attribute prediction-induced disclosure



- Train an ML model on the **released (synthetic)** data (quasi-identifiers $\rightarrow$ sensitive), then apply it to the original records.
- RAPID measures **inference risk**: how much better an adversary can predict the sensitive attribute *with* the released data than *without* it.
- Backends rf/cart/gbm/lm/logit; confidence intervals and a permutation test make it a defensible number, not a point guess.

Shipped as `rapid()` in `riskutility`; method paper under review (EPJ Data Science).

From methods to the field — and back

SNF Bridge 2023–27

longitudinal & trajectory privacy
→ riskutility

Nestlé 2026

a small, complex health study
→ pseudonymise-or-anonymise workflow

SBB 2026–

a 9M-person geo-population
→ RAPID, household swapping

Real registers, real releases, real attackers. **Each project feeds a method back into the toolbox** — the software and the consulting are one research programme.

Selected earlier work: World Bank, OECD, IHSN, Eurostat, Swisscom, Helsana.

SNF Bridge — privacy for longitudinal & trajectory data

- SNSF **Bridge Discovery** (SNSF + Innosuisse), 2023–2027. PI Templ (**FHNW**), co-PI Sariyar (BFH).
- Three fronts:
 - mask geo-location / trajectory data;
 - protect event-history & life-course data — **the sequence itself identifies**, not just the values;
 - risk assessment **beyond k -anonymity** — density-based metrics, realistic attacker simulations.
- Validated on real longitudinal data (Swiss TREE panel); tools released open (sdcMicro, riskutility).

Nestlé — anonymizing a small health study

- A nutrition study: ~ 80 children — but many sensitive variables and longitudinal follow-up. High-dimensional disclosure risk in a tiny sample: the everyday hard case.
- Method, in order:
 - write the disclosure scenario first;
 - assess identity and attribute disclosure, not k -anonymity alone;
 - anonymize where possible; optimize pseudonymization where utility requires it.
- Framed against GDPR Art. 89(1) and the EDPB anonymization guidance.

Small n is exactly where synthetic data disappoints and record swapping wins — a recurring lesson.

SBB SynPop — a 9-million-person geo-population

- Built on confidential BFS registers (STATPOP / STATENT): 9 M persons, 4 M households, 700 k businesses, at exact coordinates — barely shareable as-is.
- Two deliverables:
 - household-aware anonymization — targeted record swapping, geographic perturbation;
 - a formal disclosure-risk assessment — attribution risk, RAPID, geographic re-identification — because k -anonymity alone will not satisfy the release decision.
- Feeds transport models (SIMBA MOBi, NPVM). Partners: SBB and BFS.

Agentic AI arrives on both sides at once

AI as *assistant*

An LLM proposes anonymization strategies, scores the information loss, and refines — inside *sdcMicro*.

SoftwareX, 2026.

AI as *adversary*

A team of agents re-identifies people from “anonymised” mobility traces — cheaply, at scale, unsupervised.

PSD 2026.

Same technology — opposite ends of the table.

AI as assistant — sdcMicro's propose – evaluate – refine loop

1 Propose

LLM emits **validated tool calls** (recode, suppress, PRAM, ...)

⇒

2 Evaluate

combined **information-loss score** U

⇒

3 Refine

condition on U , propose again — until it stops improving

Only variable-level metadata is sent to the model — **never records**, not even head(data). The **practitioner signs off** the generated R code before it runs; every action leaves a reproducible script. Provider-agnostic — OpenAI, Anthropic, or a **fully on-premises** open-weight model.

AI as adversary — de-facto anonymity is a moving line

- Our PSD 2026 study: a team of AI agents re-identified **72 % of the publicly findable people** (18 of 25; 41.9 % of all 43 traces) from “anonymised” mobility data — for about **\$2 and 17 minutes** per target, unsupervised.
- “Reasonably likely means” is an **economic test**: it weighs the attacker’s cost, skill, and time. Agents collapse all three.
- So data that was *practically* anonymous can **stop being anonymous** — with no change to the data, only to the adversary. De-facto anonymity must be **re-assessed**, not assumed.
- And this is the floor, not the ceiling: agentic outreach (image geolocation, people-search, face recognition) and local open-weight models (no API audit trail) push it further.

Which is exactly your agenda

If the *empirical* anonymity line keeps moving, the answer is to stop relying on it alone:

Guarantee

Formal, provable protection — differential privacy, k -anonymity done properly — and **crypto-graphic** techniques (MPC, HE).

SDC4Society's core.

Measure

Empirical, adversarial testing — RAPID, agentic red-teaming — to know **where you actually stand**.
What we bring.

Measure the risk, then guarantee the bound. These two halves belong together.

Why FHNW × SDC4Society

- **Complementary strengths.** Your side: cryptography and formal guarantees. Our side: SDC software, empirical/adversarial risk assessment, and applied release pipelines.
- **Shared data.** Mobility and official-statistics microdata sit at the center of both agendas — SBB SynPop on our side, official statistics on yours.
- **An existing bridge.** We already share a community and reviewers — this meeting is the easy next step.

Two concrete ideas follow — a visit, and a joint project; the visit is where we draft the project.

Idea 1 — a funded research & teaching visit

The idea

A stay of a few weeks up to **~2 months**: an intensive **lecture / seminar series** on SDC, synthetic data and anonymization with R (`sdcMicro`, `simPop`/`synvey`, `riskutility`/`RAPID`), and we use the research time to **draft a joint paper and a grant proposal** (Idea 2).

- **Vehicle**: a host-applied **JSPS Invitational Fellowship** (short-term 14–60 days, or long-term 2–10 months). JSPS — not your institute's budget — covers my air travel, a living allowance, insurance and a research allowance.
- **Why a fellowship**: at FHNW every working hour is billed to a funded project, so a visit needs its own funding to be real.
- **Option**: a host-university **visiting professorship** can add formal teaching (and an honorarium) alongside the fellowship.
- Non-teaching time combined with **family travel**; timing flexible.

Idea 2 — a joint research project

A seed topic

“SDC under **agentic-AI adversaries** — formal guarantees meet empirical red-teaming for mobility & population microdata.”

Two routes fit Switzerland + Japan — at different sizes:

- **SNSF–JSPS Strategic Japanese–Swiss Programme** — bilateral, $\leq$ CHF 300k / $\leq$ JPY 30M, 24–36 months. *Right-sized for two labs*; the 2025 call has closed, we watch for the next.
- **Horizon Europe, Pillar II** — Japan fully associable since Jan 2026. Reality check: a spot-on topic (synthetic data / anonymisation) *just closed* in Apr 2026; the nearest open one (security & privacy of AI models) fits but closes Sept 2026 — too soon. A Horizon consortium also needs **more partners** (3+ countries) — so build it early and aim at the **2027 tranche**.


Thank you — and two asks

Questions welcome.

1. A **funded research & teaching visit** — as a JSPS Invitational Fellow (you apply, JSPS funds).
2. A **joint proposal** — SNSF–JSPS now, a Horizon 2027 consortium next.

matthias.templ@fhnw.ch

swissanon.com/sdc4society



Slides, packages & papers

sdcMicro, riskutility on CRAN ■ synvey on GitHub ■ SoftwareX (AI-assisted sdcMicro) ■ PSD 2026 (agentic re-identification of mobility microdata).