NAIROBI / AUDIT & EVIDENCE REFERENCE

Research perspectives

Three published discussions of connected synthetic city data.

Source and scope

This document reformats the complete published source. Historical results and current-tense wording belong to that source context. It is not a new rerun or independent certification. People, histories and model counts are synthetic.

Read before use

Persistent identities

How shared synthetic identities connect records across domains.

Historical generation

How periods and chronology shape connected life events.

Scoped checks

How published criteria, failures and limitations inform review. These discussions are not independent validation or peer-review claims.

Request the full supporting notes from Fidelra

Full published source

Nairobi · Audit report

FCTE · Evidence Library

Published audit report · Read the scope with the result

Public methodology view.

Concepts, design intent and public outcomes are described here. Internal construction methods remain controlled.

Paper 1 of 3

Connected-domain synthesis

Why a City Twin™ isn't 20 separate synthetic datasets sitting next to each other.

Every domain FCTE generates - health, employment, housing, land, insurance, and the rest - resolves back to exactly one of two canonical identities: a person (Citizen) or an organization (Business). A citizen's employer is a real row in the Business domain, not a text label; their household is a real row in the Household domain; their insurance coverage references their actual employment history, not an independent random draw. That single shared root is what makes cross-domain questions answerable at all - "does housing informality change utility-connection odds?" or "does employment status predict insurance coverage?" - because the connections are real recorded relationships, not coincidences of matching names. It's also why a change in one domain (a citizen's death, a business closing) has to propagate correctly everywhere that citizen or business is referenced, which is exactly what the platform's temporal-integrity and relationship checks are built to catch.

Paper 2 of 3

Historical reconstruction

Why the timeline starts in 1940 and is built forward, not assembled backward from today.

A synthetic "history" that's just today's snapshot with older-looking dates stamped on it doesn't hold up under scrutiny - birth cohorts wouldn't track real demographic events, employment histories wouldn't respect when institutions actually existed, and death records could reference events after a person's own death. FCTE instead generates forward from a single fixed starting point (1940) using real historical and demographic

reference data as it goes, so a citizen born in 1962 has a plausible life shaped by the events of their actual decades, not a life invented purely to match a 2026 boundary. This is also why temporal-integrity checking is a first-class, platform-wide concern rather than a one-off validation step: every domain that adds a date-bearing record is checked against the same guarantee - nothing happens before birth, nothing happens after death, nothing is ordered before its own cause.

Paper 3 of 3

Certification design

Why "certified" is six separate questions, not one score.

A single pass/fail number hides more than it reveals - a dataset can be perfectly consistent internally while resting on an unrealistic assumption, or perfectly realistic in its distributions while quietly violating referential integrity. FCTE's certification framework asks six independent questions instead: do the structures and relationships hold together (Structural Integrity), do the distributions and counts look real (Statistical Integrity), does the timeline behave (Historical Integrity), is every record genuinely synthetic (Privacy Assurance), does the same input reliably produce the same certified output (Reproducibility), and can every number be traced to a documented source or assumption (Explainability)? A release only earns its badge when all six are satisfied - and where a check currently fails, the release still ships with that failure disclosed rather than hidden, which is the same posture this Evidence Library takes.

These write-ups describe design intent and public outcomes only. Internal construction methods remain controlled - see Certification for the live pass/fail results these ideas produce.

Keep the supporting record

Ask Fidelra for the report context, applicable release, source references and full supporting notes. Availability and rights for a paid data package are confirmed in the written quote.

Request full notes and supporting artifacts

fidelra.africa@fidelra.ai

stripeddonkey.com

Original HTML source is attached for reference. Public data-quality datasets are also attached where applicable.

Source page: https://stripeddonkey.com/cities/nairobi/audits/research-papers

Captured source SHA-256: 908398416ffceb1b2ae3040a57bfd8da73322320095eae2cd3b88be838f949b0