How Eatomate Protects Your Data

Privacy Architecture — March 2026

A note from the founder

I use Eatomate every day. My family's grocery receipts, my meal logs, my body composition data — it's all in the system. The privacy architecture described in this document protects my data with exactly the same rules as yours. I built it this way because I wouldn't use a product that treated my data carelessly, and I won't ask you to either.

— Rohith, Founder & DPO, Eatomate

Our philosophy

Privacy is not a feature we added. It's how the system was designed from the first line of code.

Most companies protect your data with policies — rules that employees are told to follow. Eatomate protects your data with architecture — the system is built so that unsafe uses of your data are technically impossible, not just against the rules.

We call this **privacy by design, not by policy**.

What we collect and why

Eatomate needs certain information to give you accurate nutrition tracking. Here's exactly what we collect, and why.

To create your account:

- Name, email address, date of birth

To track your nutrition:

- What's in your pantry (items and weights)
- Your grocery receipts (items, prices, store)
- What you eat (meals, portions)
- Your exercise activity

To personalise your experience:

- Height, weight, body measurements

- Dietary preference, fitness goals, activity level

Everything we collect exists to make the product work for you. We don't collect data speculatively, and we don't collect anything we don't need.

What we don't do

This section matters more than any other.

We never sell your personal data. No advertiser, brand, or third party can buy a list of Eatomate users or their individual information. This is not a future plan — it is an architectural impossibility. Your personal identity is stored in a system that has no connection to our analytics systems.

We never link your identity to analytics data. When we produce population-level health or pricing statistics, your name, email, and personal details are not present — not anonymised, not pseudonymised, not masked. They are structurally absent. The analytics system has no way to look up who you are, because the connection doesn't exist.

We never track you at postcode level. For analytics purposes, your location is derived from the stores you shop at, rounded to borough level. We don't know your home address and we don't want to.

We never store precise timestamps. Your purchase dates are recorded at day level only. We never store the time of day you shopped, logged a meal, or exercised. This is a deliberate design choice that limits what could be inferred from your activity patterns.

We never return individual data to anyone. Every query against our analytics systems returns population averages, not individual records. There is no API endpoint, no export function, and no internal tool that returns one person's data from the analytics layer.

Two separate systems, by design

Eatomate splits data into two completely independent systems. This is the most important architectural decision in the product.

Your personal account contains everything you'd expect: your name, your meals, your pantry, your progress. This is what powers the app you use every day. It's protected by encryption, 2FA, and standard application security. Only you can see your data.

The analytics layer contains anonymous population statistics. It has no names, no emails, no dates of birth, no exact ages, no addresses. It uses anonymous identifiers that were generated randomly and cannot be traced back to your account.

The critical point: **there is no link between the two systems**. They use different identifiers that were created independently. Even if someone gained access to both systems simultaneously, there would be no way to connect a record in one to a record in the other. This isn't a policy — it's a structural property of how the identifiers were generated.

How analytics data is used

When Eatomate has enough users in a given area, we can produce population-level insights that are genuinely useful for public health. Examples:

Queries we can answer:

- "What is the average iron intake for vegetarians in Camden?"
- "Are residents of Hackney meeting the recommended daily calcium intake?"
- "How has vitamin D status changed across London boroughs this winter?"

Queries we cannot answer (by design):

- "What did the person at 14 Maple Street eat last Tuesday?"
- "Show me the diet of the vegan in postcode NW1 7BJ"
- "Which specific user is iron deficient?"

The difference isn't just that we choose not to answer the second type. The system is built so that those queries cannot execute. They are rejected before they run.

How we protect population statistics

When researchers or public health bodies query our analytics data, several protections apply automatically.

Minimum group sizes. Every statistic we return must be based on a meaningful number of people. If a query would return data based on too few individuals, the query is rejected entirely — no result is returned, and no information is leaked about how many people matched.

No exact counts. We never tell a researcher exactly how many people matched their query. Counts are reported in broad bands only. This prevents a technique where someone could subtract two similar queries to isolate one person's data.

Restricted queries. Researchers cannot combine too many filters in a single query. Asking for "the average iron intake of all vegetarians in Camden" is fine. Progressively narrowing to "vegan males, aged 30–34, in Camden, doing resistance training, with a specific body fat range" would eventually identify a single person — so the system blocks queries before they get that narrow.

Cumulative tracking. We don't just check each query in isolation. We track the total set of filters a research organisation has ever used. If their combined query history would narrow the identifiable population below our minimum threshold, their next query is blocked — even if that individual query looks harmless on its own.

If a query could identify a single person, the system simply refuses to answer.

No raw access. All access to the analytics layer — including by Eatomate employees — is mediated through controlled systems that enforce all of the above. Nobody queries the raw data directly.

Grocery pricing data

When you scan a receipt, we extract pricing information (which product, which store, what price) to build grocery intelligence — helping track price trends, detect shrinkflation, and compare value across stores.

This pricing data is separated from your identity at the moment it's created. The pricing database contains no user accounts, no names, no emails. It uses a separate anonymous identifier that cannot be connected to your personal account or to the health analytics system.

A note on health attributes in pricing data. When a purchase is recorded, a broad health snapshot is attached — for example, whether the buyer's iron status falls in the "sufficient" or "deficient" band, or their body fat percentage range (e.g. 15–20%). These are broad categories, not exact values, and they are attached to the anonymous identifier — never to your name or account. This allows aggregate queries like "what percentage of buyers of this product are iron-deficient" — useful for food brands making healthier products. No one can learn your individual health status from this data. The same minimum group size and query restrictions that protect the health analytics system apply here.

Your individual shopping basket cannot be reconstructed from the pricing database. The system prohibits any query that would combine a person's identity with both a specific date and a specific store — the only combination that could reveal what one person bought in a single trip.

Your rights

See your data. You can view everything Eatomate holds about you directly in the app.

Delete your data. You can request deletion of your account and all associated data. This includes removing your records from both the personal system and the analytics layer.

Object to analytics. You can request that your data not be included in population statistics. Contact support@eatomate.co.uk.

Ask questions. If anything in this document is unclear, email privacy@eatomate.co.uk. The founder reads every message.

Our promise

Your data is as protected as mine. The system I built for you is the system my family uses. The privacy architecture described here isn't aspirational — it's running in production, today, on my data.

If we ever change how your data is handled, we'll tell you before it happens, not after.

— Rohith

eatomate.co.uk