Compiling FIGARO: Integrating Confidentiality Treatment into Supply-Use Table Production

Johannes Buck
European Commission, Eurostat
Luxembourg

Abstract: FIGARO (Full International and Global Accounts for Research in input-Output analysis) provides annual intercountry supply, use and input-output tables for European Union (EU) Member States, EFTA and candidate countries, as well as major EU trading partners. These tables are a key resource for economic analysis, policy evaluation and international comparisons. However, the underlying data collected from national statistical institutes can contain confidential information that cannot be directly published. This paper describes the methodology developed to produce a public-use version of the FIGARO supply and use tables while protecting confidential data. The approach combines pre-masking of confidential national accounts main aggregates (NAMA) and product totals with a controlled tabular adjustment technique that is transformed into a linear programming problem by pre-determining the direction of adjustments. The method aims to minimise changes to non-confidential cells while ensuring that all published values respect confidentiality requirements and maintain the accounting identities of supply and use tables. We apply the method to all countries and years, examining the extent to which confidential cells are altered while minimising changes to non-confidential data. We analyse cases where the linear program initially fails to converge, showing how the targeted relaxation of protection levels for dominant cells resolves infeasibility.

Keywords: Supply and use tables, statistical confidentiality, controlled tabular adjustment, FIGARO, RAS method

1. Introduction

Supply and use tables (SUTs) are a fundamental tool in national accounts, providing a detailed picture of the production and use of goods and services within an economy. They record how products are supplied by domestic industries and imports, and how they are used as intermediate consumption, final consumption, investment, and exports. By construction, total supply equals total use for each product, and the system balances the production, income and expenditure approaches to gross domestic product. SUTs form the basis for compiling input-output tables, which in turn are essential for various economic analyses, including impact assessments, productivity studies, and environmental-economic accounting. At the international level, SUTs enable comparisons across countries and support the construction of multi-country supply and use and input-output models.

The FIGARO (Full International and Global Accounts for Research in input-Output analysis) process, developed by Eurostat in collaboration with the Joint Research Centre, produces inter country SUTs and input-output tables for European Union Member States, EFTA countries, candidate countries, and major EU trading partners. These tables are consistent with the European System of Accounts (ESA 2010) (Eurostat (European Commission), 2013) and follow internationally agreed classifications, including the Nomenclature of Economic Activities (NACE) for industries (Eurostat (European Commission), 2008) and the Statistical Classification of Products by Activity (CPA) for products.. FIGARO tables are widely used by researchers, policymakers, and international organisations for applications ranging from analysing global value chains to assessing economic and environmental policy impacts. Eurostat publishes derived products including macroeconomic globalisation indicators, and environmental footprints based on FIGARO. The framework also serves as the foundation for statistics under development, such as the European space economy thematic account and FIGARO national accounting matrices.

A critical challenge in publishing SUTs is the protection of confidential data. National statistical institutes provide data that may be linked to individual enterprises, which cannot be disclosed under statistical confidentiality rules. The goal of the confidentiality treatment described in this paper is to produce a dissemination version where all confidential cells have been modified to prevent disclosure, while non-confidential cells are changed as little as possible.

The approach relies on controlled tabular adjustment (CTA) with pre-masking of aggregate totals. CTA is a well-established method for protecting statistical tables by adding small perturbations to sensitive cells (Cox, 2004) (Hernández-García, 2014) (Castro, 2012). However, standard CTA involves mixed-integer programming, which becomes computationally prohibitive given the scale of FIGARO with millions of cells. To address this, we transform the problem into a linear program by pre-determining for each confidential cell whether it should be adjusted upward or downward. This decision is based on the structure of the table and feasibility considerations.

An important feature of our approach is that we first mask the national accounts main aggregates (NAMA) (output and intermediate consumption by industry) and the supply totals by product. These masked totals then serve as fixed row and column targets in the subsequent cell-level adjustment. This ensures consistency with the broader national accounts framework and reduces the dimensionality of the problem.

The paper is organised as follows. Section 2 provides an overview of the FIGARO compilation process, explaining how the tables are produced before confidentiality treatment. Section 3 describes in detail the confidentiality treatment methodology, including pre-masking of aggregates, construction of the adjustment direction matrix, feasibility checks, and the linear programming formulation. Section 4 discusses practical considerations, implementation details, and handling of special cases. It also discusses the resolution of infeasible cases. Section 5 presents the results. Finally, Section 6 provides an outlook and conclusions.

2. Compilation of FIGARO Supply and Use Tables

The compilation of FIGARO supply and use tables follows a well-established process that integrates data from multiple sources, applies harmonised methodologies, and ensures consistency across countries and over time. This section provides an overview of the main steps, focusing on aspects relevant to the subsequent confidentiality treatment.

2.1 Data Sources and Input Tables

Eurostat collects annual supply and use tables from national statistical institutes of European Union Member States, EFTA countries, and candidate countries. The data collection follows a standardised template defined in the ESA 2010 transmission programme. The main tables collected are:

- **T1500**: Supply table at basic prices, including transformation into purchasers' prices
- **T1600**: Use table at purchasers' prices
- **T1610**: Use table at basic prices
- **T1611**: Use table for domestic output at basic prices
- **T1612**: Use table for imported products at basic prices
- **T1620**: Table of trade and transport margins
- **T1630**: Table of taxes less subsidies on products

While T1500 and T1600 are collected annually, the use table in basic prices (T1610) including the detailed breakdown into domestic and imported use (T1611, T1612) and the table of trade and transport margins and taxes less subsidies (T1620, T1630) is only required every five years for years ending in 0 or 5. However, many countries provide these tables more frequently, which improves the quality of the annual estimates.

In addition to the SUT data, Eurostat collects annual national accounts aggregates (NAMA) data. These include the following by NACE Rev. 2 A*64 industry breakdown:

- **P1**: Output
- **P2**: Intermediate consumption
- **D1**: Compensation of employees
- **B2A3G**: Gross operating surplus and mixed income
- **P51C**: Consumption of fixed capital
- **D29X39**: Other taxes less subsidies on production
- **B1G**: Gross value added

Moreover, final consumption expenditure (broken down by government, households, NPISH), gross fixed capital formation, changes in inventories, exports and imports are collected. The NAMA data provide consistent time series at a more aggregate level and serve as benchmarks for the SUT compilation.

Additionally, the EU major trading partners are presented in the FIGARO tables as separate economies. For these countries, data are obtained from other international organisations, such as the Organisation for Economic Co-operation and Development (OECD), the United Nations and the World Bank, or through the websites of national statistical institutes. These data are harmonised to the extent possible with the EU classifications.

2.2 Estimation and Benchmarking

A key step in the compilation process is to ensure that the supply and use tables are consistent with the national accounts aggregates and that the various tables are balanced with each other. This is achieved through a process known as SUTRAS (Supply and Use Table Reconciliation and Adjustment System), which applies iterative proportional fitting (RAS-type methods) to adjust the tables to meet given row and column totals.

The primary data supplied by countries are stored in a database we refer to as SUT_PRIMARY. The tables are then benchmarked to national accounts aggregates. Industry-level output (P1) and intermediate consumption (P2) serve as column targets for the supply and use tables. Final use components (consumption, investment and exports) and imports are also benchmarked to the main aggregates to ensure consistency across all accounts. This ensures the SUTs reflect the most recent national accounts data, as some countries do not update their SUTs annually while the NAMA are updated every year.

For years where detailed use tables are missing, we estimate them using SUTRAS. This applies particularly to tables that are provided only every five years, i.e. to T1611 (domestic output), T1612 (imports), T1620 (trade and transport margins) and T1630 (taxes less subsidies on products). For example, when T1611 is not available for a given year but NAMA data exists, we take the structure from a base year and adjust it to the NAMA totals. The adjustment ensures consistency with output, intermediate consumption and value added. After estimation, the identity T1610 = T1611 + T1612 holds by construction.

The resulting tables are stored in a database called SUT_REFERENCE. They are fully balanced, all missing values are imputed, and all accounting identities are satisfied. However, confidential information is still present.

2.3 Extension to EU Trading Partners and Rest of the World

FIGARO covers not only EU countries but also a set of major trading partners, including the United States, China, Japan, India, Russia and others. For these countries, data are obtained from international organisations, such as the OECD, and other sources, and are adjusted to be consistent with the ESA 2010 framework where possible. A rest-of-the-world account is constructed to balance the international transactions.

To move from national SUTs to inter-country SUTs, the import in each national supply table and the exports in each national use table must be split by country of origin or destination. The exports must be additionally split by type of use (intermediate or final). This requires a detailed and balanced picture of bilateral trade flows for both goods and services.

A key challenge is that raw trade statistics often contain asymmetries – country A's reported exports to country B do not match country B's reported imports from country A. For FIGARO, a specific approach was developed to eliminate these asymmetries.

For goods, data from EU Comext and UN Comtrade are combined. The data are recoded into a unified structure, and mirror data are used to impute non-allocated trade. To resolve the remaining asymmetries, a balancing process based on a symmetry or reliability index is applied. The balanced trade flow is then broken down into three categories using the QDR methodology: quasi-transit trade (Q), domestic trade (D), and re-exports (R). This breakdown is essential, as it allows the export vector of the national use tables to be correctly allocated by partner country. Finally, adjustments are made for goods sent abroad for processing and merchanting activities to align the trade data with national accounts concepts.

For services, the process is complicated by the use of different classifications (EBOPS vs. CPA). Once a complete dataset of bilateral service flows is assembled, the same balancing approach based on a symmetry index is used to estimate a single value for each bilateral flow.

The result is a globally consistent set of trade data where global exports match global imports, valued at FOB (free on board) . This balanced trade data is then integrated with the balanced national SUTs. The final compilation step establishes the inter-country supply and use tables, as well as the derived product-by-product and industry-by-industry inter-country input-output tables. These tables cover all 27 EU Member States, two EFTA countries, five candidate countries, 15 main EU trading partners, and a 'Rest of the World' region.

2.4 From Reference to Dissemination

The SUT_REFERENCE tables are the starting point for the confidentiality treatment described in this paper. Before compiling inter-country tables by breaking down imports and exports, the goal is to produce SUT_DISSEMINATION tables that:

- Contain no confidential information (all sensitive cells have been modified)
- Are consistent with all main aggregates totals that are free for publication
- Maintain the accounting identities (e.g., row sums equal column sums, T1610 = T1611 + T1612)
- Are as close as possible to the original reference values, especially for non-confidential cells.

The following section describes this methodology in detail.

3. Confidentiality Treatment Methodology

The confidentiality treatment introduced in this paper is a multi-stage process that combines masking of aggregate totals with cell-level adjustments. The approach is designed to be computationally feasible while achieving good protection of confidential data and minimal disturbance to non-confidential data.

3.1 Overall Approach

Controlled tabular adjustment (CTA) is a statistical disclosure control method that protects sensitive cells in a table by adding small perturbations. The idea is to replace original values with adjusted ones that are sufficiently different from the true values while preserving the table's row and column totals and, where possible, keeping changes to non-sensitive cells minimal. Let $a_{i,j}$ be the original value in cell (i,j) and $x_{i,j}$ the published value. Let $\delta_{i,j}$ denote the minimum masking for cell (i,j) and $w_{i,j}$ the weight for changing the cell. For a table with n_{row} rows and n_{col} columns, the problem can be stated as:

$$\text{minimize} \sum_{i=1}^{n_{row}} \sum_{j=1}^{n_{col}} w_{i,j} |x_{i,j} - a_{i,j}|$$

$$\text{subject to} \sum_{j=1}^{n_{col}} x_{i,j} = R_i \quad \forall i$$

$$\sum_{i=1}^{n_{row}} x_{i,j} = \sum_{i=1}^{n_{row}} a_{i,j} \quad \forall j$$

$$\sum_{j=1}^{n_{col}} x_{i,j} = \sum_{j=1}^{n_{col}} a_{i,j} \quad \forall i$$

$$|x_{i,j} - a_{i,j}| \geq \delta_{i,j} \quad \text{if cell } (i,j) \text{ is confidential}$$

Standard CTA is formulated as a mixed-integer programming problem, because for each sensitive cell one must decide whether to adjust it upward or downward. While optimal, this becomes computationally prohibitive for large tables.

The approach developed for FIGARO adapts CTA to make it feasible for the scale of our production system. We transform the mixed-integer problem into a linear program by pre-determining, for each confidential cell, whether it should be increased or decreased. Once these directions are fixed, the problem becomes linear and can be solved efficiently even for large tables. The directions are chosen based on the structure of the table, the relationships between reference and primary values, and feasibility considerations.

Moreover, we also want to put additional conditions in place. For example, we demand that $x_{i,j} \geq 0$ for all cells (i,j) except when the original primary value is already negative or values which frequently turn negative, e.g. for changes in inventories. Moreover, we demand $x_{i,j} = 0$ if $a_{i,j} = 0$. This is because we might otherwise introduce values that do not make economic sense. Also, we introduce a maximum change so cells are not altered too much. Finally, the weight matrix W introduces the weight of changing a cell. It should be dependent on the size of the original value (changing a value from 10 to 20 should be more expensive than changing a value from 10,000 to 10,010) and whether the cell is flagged as confidential, assigning a larger weight for non-confidential cells.

The process consists of three main phases:

1. **Pre-masking of NAMA totals**: The output and intermediate consumption by industry (P1 and P2) as well as other NAMA values are masked to produce consistent totals that will serve as column targets for the supply and use tables.

2. **Pre-masking of supply totals by product**: If the total supply at basic prices by product is confidential (and is insufficiently different from the primary data), it is masked as well. These serve as row targets for the supply and use tables.

3. **Cell-level adjustment of detailed tables**: For each of the three main tables (T1500, T1611, T1612), and for each country and year, we solve a linear program that adjusts the cells to meet the pre-masked row and column totals while applying minimum perturbations to confidential cells and minimising changes to non-confidential cells. We do not mask T1620 and T1630 as we publish data

in basic prices. Hence, trade and transport margins and taxes less subsidies are not publicly shared. However, the data could be masked similarly.

3.2 Pre-masking of NAMA Totals

The NAMA aggregates are masked before they are used in the supply and use framework, since they act as column totals and therefore determine the overall consistency of the system. Applying the confidentiality treatment at this stage ensures that all subsequent SUT adjustments are anchored to totals that are already suitable for dissemination.

The masking is carried out at the level of industries (activities), separately for each country and year. For a given country and year, all activities flagged as confidential (or not for publication) are grouped together.

The protection mechanism is based on stochastic perturbation. For each observation, a multiplicative noise term is applied in log-space. Denoting the original value by $x_{i,t}$, the perturbed value $\tilde{x}_{i,t}$ is defined as

$$\tilde{x}_{i,t} = x_{i.t} \exp(\epsilon_{\mathrm{i,t}})$$

where $\epsilon_{i,t}$ is a zero-mean random noise term.

A key aspect of the approach is that the noise is not drawn independently over time. Instead, it is constructed to be persistent, so that the resulting series remain smooth and economically interpretable. Conceptually, the noise follows an autoregressive process of order one:

$$\epsilon_{i,t} = \rho\epsilon_{i,t-1} + u_{i,t}, \text{ where } \ u_{i,t} \sim N(\mu, \sigma^2)$$

In practice, this process is implemented via a finite moving-average representation with a sufficiently long lag length, so that shocks decay gradually over time. This avoids introducing artificial temporal volatility and helps preserve the growth patterns of the original data.

The variance is chosen to strike a balance between confidentiality protection and analytical usefulness, in particular with respect to maintaining meaningful growth rates.

To ensure reproducibility, the perturbation process is fully deterministic. For a given input dataset, the procedure always yields the same output, ensuring that revisions are not driven by the masking mechanism itself.

After applying the noise, the perturbed values are adjusted to preserve the original totals for each transaction and year. This is done by rescaling the values within each group so that their sum remains unchanged. Positive and negative values are handled separately, ensuring that both the overall total and the sign structure of the data are maintained.

Finally, the standard accounting identities are re-imposed. In particular, balancing items are recomputed from their defining components ensuring full internal consistency of the accounts.

The resulting masked NAMA aggregates are then used as fixed targets in the subsequent adjustment of the supply and use tables. For the domestic (T1611) and import (T1612) use tables, a RAS procedure is applied to allocate the adjustments in intermediate consumption between domestically produced and imported goods.

3.3 Pre-masking of Supply Totals by Product

Product totals are masked after the treatment of NAMA aggregates, as they serve as row targets in the supply-use framework. Ensuring that these totals are already confidentiality-protected guarantees consistency across all subsequent balancing steps in the tables.

The masking is applied at the level of products, separately for each country and year. For a given country–year combination, all product totals flagged as confidential (or not for publication) are treated jointly.

The protection mechanism again relies on stochastic perturbation. Multiplicative noise is applied to each observation in log-space, using the same general framework as for the activity data. The resulting perturbed values are then rescaled within each product group so that the sum of the original product totals are preserved exactly. This ensures that consistency at the aggregate level is maintained while individual confidential components are protected.

The resulting masked product totals serve as row targets for the supply table (T1500). For the use side, they are further propagated into the domestic and import use tables. Specifically, row targets for the domestic use table (T1611) are derived by subtracting imports of the product from the masked product totals, while the import use table (T1612) uses the import component directly. Imports by product (P7) are taken from the supply table, where they are already subject to the same masking and balancing procedure at the cell level.

In this way, the product-level totals remain fully consistent across the supply and use framework and align with the previously established confidentiality treatment of the underlying components.

Industries (NACE A*64) · Imports · Supply at basic prices · Taxes less subsidies · Trade and transport margins · Supply at purchasers' prices

Products (CPA 64)

NACE/CPA	A01	A02	A03	⋯	U	P7	TS_B	D21X31	OTTM	TS_O
CPA_A01	150	45	0	⋯	0	30	250	20	8	278
CPA_A02	10	40	5	⋯	0	20	95	5	10	110
CPA_A03	0	10	10	⋯	0	5	45	3	5	53
⋮	⋮	⋮	⋮	⋱	⋱	⋮	⋮	⋮	⋮	⋮
CPA_U	0	0	0	⋯	0	0	0	0	0	0
NAMA	165	100	20	⋯	0	300	1500	150	80	1230

Masked product totals serve as row targets

Masked industry totals serve as column targets · NAMA adjusted for CIFFOB, DPA

Table 1: Illustration of how the masked NAMA values serve as column totals and masked product totals serve as row totals in T1500.

3.4 Cell-Level Adjustment: Problem Formulation

For each table, country and year we solve the same core problem. Let $x_{i,j}$ be the published value for product i and column j. The reference value $r_{i,j}$ is our starting point.

We have row targets R_i from masked supply totals and column targets C_j from masked NAMA totals. Moreover, let $p_{i,j}$ be the primary value transmitted by the country.

The reference values correspond to the balanced supply and use tables benchmarked to national accounts main aggregates (NAMA), while the primary data reflect the original reporting from national statistical institutes before any balancing or benchmarking adjustments.

The problem can be stated as:

$$\text{minimize} \sum_{i=1}^{n_{row}} \sum_{j=1}^{n_{col}} w_{i,j} \left| x_{i,j} - r_{i,j} \right|$$

$$\text{subject to} \sum_{j=1}^{n_{col}} x_{i,j} = R_i \quad \forall i$$

$$\sum_{i=1}^{n_{row}} x_{i,j} = C_j \quad \forall j$$

$$L_{i,j} \in \{-1,1\} \text{ if cell } (i,j) \text{ is confidential}$$

$$L_{i,j} = 0 \text{ if cell } (i,j) \text{ is not confidential}$$

$$x_{i,j} - p_{i,j} \geq \delta_{i,j} \quad \text{if cell } (i,j) \text{ is confidential and } L_{i,j} = 1$$

$$p_{i,j} - x_{i,j} \geq \delta_{i,j} \quad \text{if cell } (i,j) \text{ is confidential and } L_{i,j} = -1$$

$$x_{i,j} - r_{i,j} \leq \pi_{i,j} \quad \text{for all cells } (i,j)$$

$$r_{i,j} - x_{i,j} \leq \pi_{i,j} \quad \text{for all cells } (i,j)$$

$$x_{i,j} \geq 0 \quad \text{if } p_{i,j} \geq 0 \quad \text{and column is not P5M}$$

$$x_{i,j} = 0 \quad \text{if } p_{i,j} = 0$$

The first two constraints ensure row and column totals match the masked aggregates. Non-negativity is enforced for most cells, with changes in inventories (P5M) as the only exception. Cells that are zero in the primary data remain zero in the published table to not introduce values that do not bear any economic meaning.

For confidential cells, the published value must differ from the primary value by at least $\delta_{i,j}$, a pre-defined minimum perturbation. Finally, no cell may change by more than $\pi_{i,j}$, a bound derived from the overall discrepancy between current totals and targets.

The original formulation is a mixed-integer program because for each confidential cell we would need to decide whether $x_{i,j}$ lies above or below $p_{i,j}$. To make the problem tractable at scale, we pre-determine the direction of adjustment for every confidential cell. To do so, we set $L_{i,j} = 1$ if the cell should increase, $L_{i,j} = -1$ if it should decrease, and $L_{i,j} = 0$ for non-confidential cells. With directions fixed, the protection constraints become linear:

$$x_{i,j} - p_{i,j} \geq \delta_{i,j} \quad \text{if } L_{i,j} = 1$$

$$p_{i,j} - x_{i,j} \geq \delta_{i,j} \quad \text{if } L_{i,j} = -1$$

The absolute value in the objective can be easily linearised by introducing two non-negative variables for each cell: $u_{i,j}$ representing an increase from the reference value, and $d_{i,j}$ representing a decrease. The final value is:

$$x_{i,j} = r_{i,j} + u_{i,j} - d_{i,j}.$$

3.5 Determining Adjustment Directions

Standard controlled tabular adjustment requires deciding for each confidential cell whether it should be adjusted upward or downward. This turns the problem into a mixed-integer program, which becomes computationally infeasible at the scale of

FIGARO. We therefore determine the adjustment directions in advance and subsequently solve only the resulting linear program.

The direction assignment is constructed separately from the row and column perspectives. Within each row and column, confidential cells are ordered by size and assigned alternating directions, producing an initial pattern that is broadly balanced across the table. Where the row- and column-based assignments differ, the final direction is chosen so as to improve the overall balance of the system.

The initial assignment is then refined using the masked row and column targets. Intuitively, if the required adjustment implied by the targets points consistently in the opposite direction of the current assignment, the sign is reversed. This helps align the direction matrix with the structure of the balancing problem and improves feasibility of the subsequent optimisation.

3.6 Feasibility Checking and Direction Refinement

Before solving, we verify that the chosen directions can work. For a given direction matrix L, for each row i, we can calculate the maximum and minimum row values given the directions of confidential cells and the maximum change $\pi_{i,j}$.

Summing up the maximum and minimum values, gives us a range for the row target. Only if the true row target lies within this range, the optimization problem can be feasible.

Similar bounds exist for columns. If any target R_i lies outside its row bounds or any C_j lies outside its column bounds, the current directions cannot yield a feasible solution. We then apply a correction method to resolve the infeasibility.

3.7 Summary of the Algorithm

We are now ready to state the algorithm formally. The procedure combines the pre-masking steps with the cell-level linear program into a coherent workflow that runs independently for each country, year and table.

The algorithm proceeds as follows:

1. **Pre-mask NAMA totals:** Confidential NAMA aggregates are masked at the level of industries, separately for each country and year. The protection mechanism is based on stochastic perturbation with persistent noise over time, followed by a rescaling step to preserve the original totals for each transaction. Accounting identities are then re-imposed by recalculating balancing items from their defining components. The resulting masked values are stored as

column targets C_j. For T1611 and T1612, a RAS procedure is subsequently used to distribute the adjustments in intermediate consumption between domestic and imported use.

2. **Pre-mask supply totals by product:** Confidential product totals are masked separately for each country and year using the same stochastic perturbation framework as for the NAMA aggregates. After perturbation, values are rescaled so that the overall product totals remain unchanged. The resulting masked values are stored as row targets R_i. For T1600, margins and taxes are added to obtain purchasers' price targets. For T1611 and T1612, the domestic use targets are derived by subtracting imports from the masked supply totals, while the import use targets are taken directly from the masked imports in the supply table.

3. **For each table (T1500, T1611, T1612):**

   a. **Construct the direction matrix $\boldsymbol{L}$.**

   b. **Check feasibility.**

   c. **Solve the linear program.**

   d. **If the linear program with current directions is infeasible**, apply fallback (see details below).

   e. **Store the adjusted cells** Update the dissemination database with the new values.

4. **Reconstruct derived aggregates** After all tables are processed, recalculate T1610 as the sum of T1611 and T1612 and T1600 as the sum of T1610, T1620 and T1630. After all tables are processed, recalculate all row and column totals, including supply totals, final use aggregates and overall use totals, from the adjusted detailed cells to ensure accounting consistency.

The algorithm is implemented in R using the *lpSolve* package for linear programming. Each country-year-table combination is processed independently, which allows for parallelisation and keeps memory requirements manageable. The feasibility correction loop typically converges within a few iterations, and the linear program solves in seconds for tables of typical size.

A pseudo-code representation of the procedure is provided in Algorithm 1 below.

Algorithm 1 Controlled Tabular Adjustment for Supply and Use Tables

Require: Country c, table $t \in \{T1500, T1611, T1612\}$ and year y
Require: Reference matrix M and primary matrix P
Require: Row totals $R_i,\ i = 1, \ldots, n_{\text{row}}$
Require: Column totals $C_j,\ j = 1, \ldots, n_{\text{col}}$
Require: Weight matrix W

1: Load confidential cells and corresponding primary values from P
2: Apply pre-treatment:
  identify structural zero cells, non-sensitive cells and fixed cells
  define admissible bounds for cell changes
3: Construct minimum protection levels δ_{ij}
  based on confidentiality status and dominance criteria relative to P
4: Construct maximum admissible changes π_{ij} relative to M
  based on discrepancies between current and target row and column totals
5: Construct direction matrix L
  using row and column constraints with protection requirements for confidential cells
6: Check feasibility of direction matrix L
  under protection constraints δ_{ij}, maximum changes π_{ij} and row and column targets
7: **if** feasibility conditions are violated **then**
8:   Correct directions
9: **end if**
10: Construct linear program:
  objective minimizes weighted cell adjustments relative to M and weights W
  row and column totals match their targets
  protection levels δ_{ij} must hold
  maximum changes π_{ij} are enforced
  non-negativity constraints are imposed where applicable
11: Solve linear program using `lpSolve`
12: **if** no feasible solution is found **then**
13:   Apply fallback strategy and resolve
14: **end if**
15: Compute

$$T1610 = T1611 + T1612$$

16: Compute

$$T1600 = T1610 + T1620 + T1630$$

17: Recalculate derived aggregates and higher-level totals
18: Store adjusted table and update dissemination database

Algorithm 1: Pseudo-algorithm for Controlled Tabular Adjustment with predetermined directions as part of the FIGARO confidentiality treatment.

4. Implementation, Special Cases, and Convergence

4.1 Handling Special Situations

This section discusses several special situations and how we treat them.

1. **Confidential zero values:**
   In some cases, entire rows or columns are marked as confidential or not for publication, even though they contain individual zero values. This occurs particularly in supply tables. Masking such entries would require replacing structural zeros with positive values, which can be economically implausible and introduces artificial production or flow.
   To avoid this, we do not apply masking to zero-valued cells. These values remain fixed at zero throughout the adjustment process.
2. **Confidential cells dominating the row:**
   A second special case arises when a confidential cell dominates its corresponding row or column. This is especially relevant in the supply table T1500. For example, consider an industry total of 100,000 that is free for publication, split between two products with values 99,990 and 10. If the large cell is subject to masking, even a small relative adjustment implies that the balancing change must be absorbed by the remaining small cell in order to preserve the accounting identity.
   In this situation, the small residual value would experience extremely large proportional changes, potentially in the order of several thousand percent. Such distortions are not consistent with the maximum change constraints in the optimization problem and would lead to economically implausible reallocations.
   To address this issue, such dominance structures are identified ex ante. In these cases, masking intensity is reduced or selectively restricted for dominant cells, ensuring that the adjustment is absorbed in a way that avoids excessive distortion of smaller components while preserving overall consistency of totals.
3. **Infeasible or borderline LP solutions:**
   In rare cases, the linear program may not admit a feasible solution. This can occur even without pre-specifying adjustment directions.
   Such situations are generally isolated and only arise for a small number of country, table and year combinations where the combination of confidentiality patterns, target constraints, and admissible adjustment bounds makes the system overly constrained.
   If infeasibility persists after addressing the special cases described above and after progressively relaxing the maximum allowable percentage change, the model may still fail to converge. In these situations, we apply a dedicated

convergence enforcement procedure, described in the following subsection. This procedure guarantees that a feasible and internally consistent solution can always be obtained.

4.2 Resolution of Infeasible Cases

When the optimization problem remains infeasible after all previous adjustments, we apply a structured feasibility enforcement procedure. The strategy differs slightly between supply and use tables but follows the same general principle of progressively relaxing constraints in a controlled manner.

Supply table (T1500)

In the supply table, infeasibility is addressed by relaxing the row targets while keeping the column targets from national accounts main aggregates fixed. This reduces the problem to matching column totals only, which is typically easier to satisfy and can be resolved using the adjustment mechanisms described in the previous subsection.

It is worth noting that in standard RAS-type balancing procedures, row totals are also treated as flexible. By construction, the sum of row totals equals the sum of column totals. If the sum of column totals is adjusted (for example through benchmarking to external NAMA P1 data), the sum of row totals must adjust accordingly to preserve this identity and hence row totals cannot be preserved.

In our setting, relaxing row constraints is not strictly necessary. Because the masking procedure preserves aggregate totals, the sum of column targets remains unchanged, and consistency could still be achieved while keeping row totals fixed.

Hence, row flexibility is introduced as a safeguard only in infeasible cases. Nevertheless, this is consistent with standard RAS-type balancing procedures, where row totals are routinely allowed to adjust to ensure reconciliation. Therefore, although not required by construction, this relaxation is methodologically acceptable and is not considered problematic.

Use tables (T1611 and T1612)

For the use tables, the situation is slightly different, as the primary requirement is to satisfy column totals, while the discrepancies are now absorbed through adjustments in inventories.

In case of infeasibility, we first relax restrictions related to the changes in inventories for goods, since these entries can take both positive and negative values and therefore allow economically meaningful flexibility. This often resolves most remaining inconsistencies.

If infeasibility persists, the restrictions on inventories are further relaxed across all rows. This may introduce non-zero adjustments for service categories where values are typically expected to be zero. However, due to the structure of the weighting scheme, such adjustments are heavily penalised, so that only minimal changes are introduced, just enough to restore feasibility.

5. Results

We now present the main empirical results of the proposed confidentiality treatment. Since the methodology is applied to real confidential FIGARO data intended for official dissemination, certain implementation details, particularly the exact minimum perturbation thresholds and related protection parameters, cannot be disclosed. Consequently, some results are intentionally reported in relative rather than absolute terms.

To preserve anonymity, country identifiers are masked throughout the analysis. All countries containing confidential or not-for-publication NAMA aggregates, or confidential or not-for-publication cells within the supply and use tables, are randomly relabelled as C01–C19, with the exception of two countries. These two countries were excluded because confidentiality was only present in the supply table, making identification substantially easier. The assigned labels are purely random and do not reflect geography, economic size, alphabetical order, or any other country characteristic.

It is important to note that part of the observed perturbation in the supply and use tables originates from prior masking of the NAMA aggregates described in Section 3.2. Since the masked NAMA values serve as column constraints within the balancing framework, changes introduced at the aggregate level necessarily propagate into the underlying SUT cells, including cells that are themselves free for publication. The results presented below therefore reflect the combined effect of aggregate-level and cell-level confidentiality treatment within the integrated balancing framework.

5.1 Preservation of Free-for-Publication Data

We first evaluate the extent to which cells that are free for publication remain unchanged after confidentiality treatment. Table 2 reports, for each country and table, the share of free-for-publication cells for which the absolute difference between SUT_REFERENCE and SUT_DISSEMINATION is smaller than 0.001 million euro. The threshold of 0.001 is used because FIGARO data are disseminated with three decimal places.

Overall, the results indicate that the overwhelming majority of observations that are free for publication remain effectively unchanged. Across all countries and tables, the median share of unchanged cells equals 98.18%, demonstrating that the methodology preserves the original structure of the benchmarked tables to a very high degree.

The lowest observed value occurs for country C04 in table T1611, where 94.89% of free-for-publication cells remain unchanged. Even in this comparatively difficult case, the vast majority of publicly available information is preserved. This confirms that the optimization framework succeeds in concentrating adjustments primarily on confidential observations while minimizing spillovers to non-confidential cells.

5.2 Changes to Confidential and Not-for-Publication Cells

We next analyse the magnitude of revisions applied to observations flagged as confidential or not for publication. Throughout this section, the revisions are measured as the difference between the balanced benchmark tables (SUT_REFERENCE) and the final disseminated tables (SUT_DISSEMINATION). Hence, the analysis captures the total effect of the confidentiality treatment after benchmarking and balancing have already been completed.

Because the interpretation of revisions depends strongly on the scale of the underlying observations, we distinguish between three value classes: small observations (below 10 million euro), medium-sized observations (between 10 and 100 million euro), and large observations (above 100 million euro).

Small Values (< 10 million euro)

For very small observations, percentage-based revision measures can be misleading. Even economically negligible absolute changes may generate extremely large percentage changes when the original value is close to zero.

For this reason, small observations are evaluated using absolute rather than relative changes. Moreover, observations equal to exactly zero are excluded from the analysis. As discussed earlier, structural zeros are generally preserved by the algorithm and including them would mechanically inflate the stability indicators.

Table 3 reports, for confidential and not-for-publication observations below 10 million euro, the percentage of cells for which the absolute difference between

SUT_REFERENCE and SUT_DISSEMINATION remains below 1, 3, and 5 million euro, respectively.

Overall, the results indicate that distortions among small observations remain limited. Even under the strictest threshold of an absolute revision below 1 million euro, the median share of cells satisfying this condition equals 90.81%. This suggests that the vast majority of small-valued observations experience only very moderate absolute adjustments during the confidentiality treatment. As the threshold is relaxed to 3 and 5 million euro, the shares increase substantially across almost all country-table combinations. For the 5 million euro threshold, the median share reaches 96.96%, indicating that only a very small proportion of small-valued observations are subject to larger absolute revisions.

The weakest results are observed particularly for countries C07 and C13. Nevertheless, even in these cases, the overwhelming majority of observations remain within comparatively moderate absolute revision ranges. This suggests that although small values are naturally more sensitive to balancing adjustments, the confidentiality treatment generally avoids economically implausible distortions.

Medium-Sized Values (10–100 million euro)

For medium-sized observations, relative changes become more meaningful and economically interpretable. However, because the exact minimum perturbation requirements used internally for confidentiality protection cannot be disclosed, the results are presented relative to the minimum required perturbation rather than in absolute percentage terms.

More precisely, for each confidential or not-for-publication observation, we calculate the ratio

$$\frac{|\mathrm{SUT\_DISSEMINATION} - \mathrm{SUT\_REFERENCE}|}{\mathrm{MINIMUM\_PERTURBATION}}$$

where the numerator corresponds to the absolute change between SUT_REFERENCE and SUT_DISSEMINATION, and the denominator corresponds to the internally imposed minimum perturbation requirement for the specific cell (in absolute terms). Table 3 reports the mean, the 75th percentile, and the 95th percentile of this ratio by country and table.

Overall, the results indicate that revisions for medium-sized observations generally remain within a moderate range. Across all countries and tables, the typical mean ratio is approximately 2.7, while the typical 75th percentile is around 3. This implies that for most medium-sized confidential observations, revisions remain limited to a multiple of only a few times the minimum required perturbation.

A few countries nevertheless exhibit mean or upper-quantile ratios below 1 for some tables, particularly C04 and C09. At first sight, this could appear inconsistent with the imposed confidentiality requirements. However, this interpretation would not be correct because the confidentiality constraints are defined relative to SUT_PRIMARY, whereas the revisions analysed here are measured between SUT_REFERENCE and SUT_DISSEMINATION. Consequently, if the benchmarking process from SUT_PRIMARY to SUT_REFERENCE already introduced significant adjustments, only limited additional perturbation may be required during the confidentiality treatment itself. In such cases, ratios below 1 can naturally occur even though all confidentiality constraints remain fully satisfied relative to the original primary data.

The results also reveal substantial differences across countries and tables. Larger revisions are observed in all three table types, although the strongest upper-tail outliers occur most frequently in T1612 and, to a lesser extent, in T1500. Particularly, large revisions are observed for countries C03, C07, C11 and C19. For example, the 95th percentile exceeds 25 for C03 in T1612. These findings indicate that while the majority of observations remain comparatively stable, a limited number of cells may require substantially larger adjustments in more constrained confidentiality settings.

The 95th percentile is particularly informative in this regard, as it captures more extreme revision scenarios rather than typical behaviour. The typical 95th percentile across all countries and tables equals approximately 9.6, indicating that the more substantial revisions involve about 10 times the minimum required perturbation. Since the 95th percentile specifically captures more extreme scenarios, these results suggest that while most observations remain comparatively stable, a limited number of cells can require considerably larger adjustments.

At the same time, the dispersion across countries demonstrates that the feasibility of the confidentiality treatment depends strongly on the structural characteristics of the underlying tables and the associated confidentiality patterns. In cases where confidentiality requirements are especially restrictive, for example because a large share of a table is confidential, the optimisation problem becomes considerably more rigid and a few larger outliers may emerge.

Despite these differences, the overall results suggest that the proposed methodology succeeds in protecting confidential information while preserving most economically relevant structures of the benchmarked supply and use tables.

Large Values (> 100 million euro)

For large observations, Table 3 applies the same ratio-based methodology as for medium-sized observations.

Compared with medium-sized observations, the results show smaller perturbations. Across all countries and tables, the typical mean ratio falls to approximately 1.2 while the typical 75th percentile equals around 1.3. This indicates that large confidential observations generally remain close to their benchmarked SUT_REFERENCE values despite the imposed confidentiality constraints.

The same structural patterns remain visible, although much less pronounced. Ratios below 1 again occur for some countries and tables, reflecting the fact that part of the required perturbation may already have been introduced during the benchmarking step from SUT_PRIMARY to SUT_REFERENCE. However, compared with medium-sized observations, the dispersion is considerably smaller and the overall stability is much higher.

The typical 95th percentile across all countries and tables equals approximately 3.8 which is substantially lower than for medium-sized observations. This suggests that even relatively extreme revisions for large observations generally remain moderate. Nevertheless, some heterogeneity across countries remains visible. In particular, larger upper-tail revisions are observed for C03, C11 and C14, whereas several other countries remain close to 1 across all tables.

Overall, the results indicate that the proposed methodology performs particularly well for economically important large-valued observations. Even under relatively restrictive confidentiality requirements, the vast majority of large cells remain close to their benchmarked values, while only a limited number of observations experience more substantial revisions in highly constrained settings.

5.3 Reduced Masking for Dominant Confidential Cells

We next investigate how often the standard perturbation requirement had to be relaxed because a confidential cell dominated its corresponding column, as discussed in Section 4.1.

Table 4 reports the share of confidential or not-for-publication cells for which a reduced masking intensity relative to SUT_PRIMARY was applied in order to preserve feasibility and avoid implausible reallocations.

Overall, such situations are relatively rare. Across all countries and tables, the median share equals approximately 4.13%. This indicates that the standard protection mechanism remains feasible in the vast majority of cases.

The phenomenon occurs considerably more frequently in the supply table T1500, where the median reaches 6.85%, compared with 3.34% and 3.69% for T1611 and T1612 respectively. This result is economically intuitive. Supply tables are typically more concentrated around the diagonal, meaning that production of a product is often dominated by a single industry. In such situations, perturbing a dominant confidential cell while preserving row and column identities can otherwise force implausibly large adjustments onto small residual cells forcing a reduced masking.

The results therefore confirm the importance of the dominance-handling mechanism introduced in Section 4.1.

5.4 Relaxation of Row Targets in the Supply Table

We next analyse cases where the optimization problem remained infeasible and row targets in the supply table T1500 had to be relaxed, as described in Section 4.2.

Table 5 reports the share of row totals for which the original SUT_REFERENCE target could not be maintained exactly and therefore had to be opened during the balancing procedure.

Overall, such cases are uncommon. The median share equals only 1.98%, indicating that the original row targets can generally be preserved even under the confidentiality constraints.

However, two strong outliers are observed for countries C04 and C19. These cases appear to arise from particularly restrictive confidentiality structures. For example, infeasibility may occur when nearly all observations in a row or column are confidential or not for publication, leaving insufficient flexibility for the optimization problem.

Another problematic configuration occurs when the supply table approaches a near-perfect diagonal structure. In this situation, a single non-zero observation may simultaneously determine both a row total and a column total. If the corresponding NAMA aggregate is confidential and therefore perturbed, the original product total can no longer remain unchanged while structural zeros are preserved. In such cases, relaxing the row target becomes necessary to restore feasibility.

Importantly, allowing row totals to adjust is fully consistent with standard RAS-type balancing procedures and is therefore not considered methodologically problematic.

5.5 Changes to Structural Zeros in Changes in Inventories (P5M)

Finally, we examine cases where structural zero values in the P5M column (changes in inventories) were converted into non-zero observations during the balancing process.

As discussed in Section 4.2, relaxing the P5M restrictions may become necessary in rare infeasible cases because inventories provide one of the few economically meaningful sources of additional flexibility within the use tables.

Table 6 reports the share of all P5M cells for which a zero value in SUT_REFERENCE became a non-zero value in SUT_DISSEMINATION.

The results are highly reassuring. The median share equals zero, implying that for most country-table combinations no structural zeros in inventories had to be modified at all.

The highest observed value occurs for country C12 in T1612, where 1.15% of P5M cells changed from zero to non-zero values. Even this value remains extremely small.

Moreover, these observations are associated with very high penalties in the optimization objective due to their size. Consequently, even when modifications occur, the resulting perturbations are typically economically negligible and remain concentrated in very small absolute magnitudes.

Table	C01	C02	C03	C04	C05	C06	C07	C08	C09	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19
T1500	100	98.01	97.66	98.88	98.61	100	98.27	100	99.09	98.34	97.19	98.41	98.59	98.06	97.85	99.99	97.72	99.89	98.52
T1611	100	97.72	97.74	94.89	97.67	100	98.25	100	99.50	98.02	97.53	97.51	97.72	98.19	97.78	100.00	98.43	99.80	97.26
T1612	100	98.04	97.51	97.65	98.04	100	98.18	100	99.19	98.11	97.52	97.66	97.49	97.25	98.17	100.00	97.48	99.84	97.87

Table 2: Share of free-for-publication cells remaining unchanged after confidentiality treatment (in %, absolute difference < 0.001)

Size	Table	Revision Indicator	C02	C03	C04	C05	C07	C09	C10	C11	C13	C14	C15	C19
	T1500	<1	90.29	91.03	99.44	84.80	72.41	100.00	72.35	89.92	74.39	90.68	90.47	96.10
	T1500	<3	93.74	95.03	99.78	88.00	74.14	100.00	78.80	92.76	85.98	93.66	94.41	98.52
	T1500	<5	95.28	97.28	99.79	90.40	74.14	100.00	83.41	94.36	90.85	94.41	95.97	99.34
	T1611	<1	95.34	97.69	95.85	77.78	92.56	99.32	93.46	90.27	72.09	85.58	89.84	98.20
small	T1611	<3	98.10	99.26	99.93	88.27	95.84	99.61	96.39	93.21	83.14	89.74	95.52	99.19
	T1611	<5	98.87	99.66	99.93	91.36	96.06	99.61	97.29	95.03	88.95	93.06	97.26	99.49
	T1612	<1	97.05	97.47	96.72	88.59	79.26	99.12	90.92	93.43	79.90	85.71	90.70	98.97
	T1612	<3	98.99	99.30	99.76	92.62	82.47	99.91	96.26	95.76	85.78	90.79	96.17	99.70
	T1612	<5	99.49	99.62	99.95	95.30	83.21	99.94	98.20	96.66	89.22	94.20	97.29	99.87
	T1500	Mean	2.77	2.90	1.28	3.32	5.84	0.94	1.59	4.35	3.70	3.65	2.53	1.88
	T1500	75%	1.71	2.78	1.00	2.62	11.46	1.00	1.25	6.06	4.02	2.57	1.62	1.19
	T1500	95%	12.28	12.89	1.26	12.23	19.62	1.07	3.50	15.12	13.17	14.22	10.97	7.98
	T1611	Mean	2.26	3.31	0.56	3.10	2.20	0.70	2.72	3.14	2.39	2.59	2.02	4.68
medium	T1611	75%	2.81	3.30	0.92	3.00	2.13	1.14	3.21	4.25	3.34	3.84	2.99	4.93
	T1611	95%	6.85	17.43	2.39	14.76	8.39	2.81	8.15	9.67	7.80	8.39	7.93	18.45
	T1612	Mean	2.49	6.83	0.91	1.59	4.03	0.83	2.97	4.55	3.20	2.85	2.37	2.75
	T1612	75%	3.65	9.04	1.46	2.97	5.97	1.29	4.25	7.23	4.50	4.40	3.16	3.21
	T1612	95%	7.41	25.98	3.13	5.90	12.05	3.67	7.25	13.70	10.64	9.72	9.60	10.13
	T1500	Mean	0.51	0.74	1.03	0.92	2.09	0.72	0.78	1.70	2.02	1.74	1.25	0.94
	T1500	75%	0.76	1.07	0.88	0.89	1.62	1.02	1.00	1.27	1.52	1.44	1.04	1.25
large	T1500	95%	1.19	2.58	2.90	5.18	11.25	1.47	1.00	9.55	9.54	7.88	9.35	2.33
	T1611	Mean	1.19	4.04	0.74	0.72	1.05	1.26	1.06	1.92	1.25	0.82	1.02	1.68
	T1611	75%	1.44	2.70	0.83	1.23	1.20	1.56	1.04	3.01	1.48	1.09	1.28	2.25
	T1611	95%	1.56	21.46	3.38	1.44	1.38	5.26	1.16	3.99	4.08	1.46	3.75	3.37

Size	Table	Revision Indicator	C02	C03	C04	C05	C07	C09	C10	C11	C13	C14	C15	C19
	T1612	Mean	0.65	2.34	1.21	1.19	1.59	1.13	1.00	2.13	1.35	1.81	1.57	1.01
	T1612	75%	0.80	2.10	1.34	1.08	1.24	1.58	1.01	3.34	1.53	1.39	1.68	1.59
	T1612	95%	1.69	4.81	5.61	3.76	4.95	4.45	1.02	9.09	3.48	8.03	4.14	2.47

Table 3: Distribution of revisions for confidential and not-for-publication cells by value class

Table	C02	C03	C04	C05	C07	C09	C10	C11	C13	C14	C15	C19
T1500	10.52	1.23	11.82	6.12	1.08	3.91	1.21	11.12	13.93	9.26	7.59	3.13
T1611	10.52	3.45	0.59	0.00	4.36	0.01	11.55	9.45	0.00	2.58	3.24	13.31
T1612	7.17	2.78	4.60	0.37	1.93	1.06	9.61	10.93	0.31	5.56	2.05	10.14

Table 4: Share of confidential or not-for-publication cells subject to reduced masking due to dominance structures (in %)

C01	C02	C03	C04	C05	C06	C07	C08	C09	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19
0.21	1.98	4.58	90.94	25.1	0	1.98	0.1	0	0.73	4.9	12.19	1.15	0	10.83	0.42	0	6.77	80.62

Table 5: Share of cases requiring relaxation of row targets in the supply table (in %)

Table	C01	C02	C03	C04	C05	C06	C07	C08	C09	C10	C11	C12	C13	C14	C15	C16	C17	C18	C19
T1611	0	0	0.21	0.00	0.00	0	0	0	0.00	0	0.0	0.00	0.00	0.00	0.00	0	0.0	0	0.00
T1612	0	0	0.10	0.31	0.83	0	0	0	0.42	0	0.1	1.15	1.04	0.21	0.21	0	0.1	0	0.62

Table 6: Share of changes in inventories (P5M) cells altered from zero to non-zero values (in %)

6. Outlook and Conclusions

The methodology described in this paper has been developed for the public-use version of the FIGARO supply, use and input-output tables. It aims to protect confidential data while keeping the tables useful for economic analysis. By combining pre-masking of aggregates with cell-level adjustments, the developed linear program runs fast enough for production and provides a solution that changes non-confidential data as little as possible.

Nevertheless, several extensions and methodological refinements could be explored in future research. One possibility would be to integrate the confidentiality treatment directly into the SUTRAS balancing step, rather than doing it afterwards. However, such an approach would also require maintaining a separate internal version of the FIGARO tables without confidentiality treatment for downstream production processes and analytical applications. As a result, the overall production workflow would become considerably more complex and computationally demanding. One could also go back to the original mixed-integer formulation and try to solve it with branch-and-bound or similar methods. This would however be computationally demanding given that we have thousands of cells per table and well over a thousand country-year-table combinations to process.

Feedback from researchers and analysts using the FIGARO disseminated tables, will also help us improve the method over time. The FIGARO methodology is still evolving, and the confidentiality treatment we have developed makes it possible to publish the FIGARO tables while respecting the confidentiality requirements of the statistical institutes that provide the underlying data.

Bibliography

Castro, J. (2012). On assessing the disclosure risk of controlled adjustment methods for statistical tabular data. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 921-941.

Cox, L. a. (2004). A Disclosure Limitation Method for Tabular Data That Preserves Accuracy and Ease-of-Use. *Proceedings of the 2002 FCSM Statistical Policy Conference, Washington, DC: Office of Management and Budget*, (pp. 15-30).

Eurostat (European Commission). (2008). *NACE Rev. 2 - Statistical classification of economic activities.* Luxembourg: Publications Office of the European Union.

Eurostat (European Commission). (2013). *European System of Accounts – ESA 2010.* Luxembourg: Publications Office of the European Union.

Eurostat (European Commission). (n.d.). *FIGARO project website*. Retrieved from Eurostat: https://ec.europa.eu/eurostat/web/esa-supply-use-input-tables/information-data#figaro

Hernández-García, M.-S. a.-G.-J. (2014). Enhanced controlled tabular adjustment. *Computers & Operations Research*, 43: 61-67.