CODEX PLANETARIUS

Codex Planetarius Pilot Study: Palm in Paradise - A Supply Shed Analysis of Indonesian Palm Oil

Epoch

About *Codex Planetarius*

Codex Planetarius is a proposed system of minimum environmental performance standards for producing globally traded food. It is modeled on the *Codex Alimentarius*, a set of minimum mandatory health and safety standards for globally traded food. The goal of *Codex Planetarius* is to measure and manage the key environmental impacts of food production, acknowledging that while some resources may be renewable, they may be consumed at a faster rate than the planet can renew them.

The global production of food has had the largest impact of any human activity on the planet. Continuing increases in population and per capita income, accompanied by dietary shifts, are putting even more pressure on the planet and its ability to regenerate renewable resources. We need to reduce food production's key impacts.

The impacts of food production are not spread evenly among producers. Data across commodities suggest that the bottom 10-20% of producers account for 60-80% of the impacts associated globally with producing any commodity, even though they produce only 5-10% of the product. We need to focus on the bottom.

Once approved, *Codex Planetarius* will provide governments and trade authorities with a baseline for environmental performance in the global trade of food and soft commodities. It won't replace what governments already do. Rather, it will help build consensus about key impacts, how to measure them, and what minimum acceptable performance should be for global trade. We need a common escalator of continuous improvement.

These papers are part of a multiyear proof of concept to answer questions and explore issues, launch an informed discussion, and help create a pathway to assess the overall viability of *Codex Planetarius*. We believe *Codex Planetarius* would improve food production and reduce its environmental impact on the planet.

This proof-of-concept research and analysis is funded by the Gordon and Betty Moore Foundation and led by World Wildlife Fund in collaboration with a number of global organizations and experts.

For more information, visit www.codexplanetarius.org

Codex Planetarius Pilot Study: Palm in Paradise - A Supply Shed Analysis of Indonesian Palm Oil

Epoch

1 Executive Summary

The *Codex Planetarius* Pilot Study, "Palm in Paradise - A Supply Shed Analysis of Indonesian Palm Oil," aims to establish a minimum performance standard for environmental impacts in the global food and soft commodity trade. This framework focuses on building consensus for measuring key impacts like habitat loss, biodiversity, and emissions across supply chains, specifically targeting the bottom 10-20% of producers responsible for the majority of environmental harm.

Epoch has contributed by proposing a methodology for first-mile mapping of commodity supply chains, from plot to first aggregation point, using a state-of-the-art supply shed delineation approach. This method establishes a probabilistic relationship between commodity-producing areas and their associated processing facilities, allowing for systematic and granular environmental performance assessment.

The pilot study applied this approach to the Indonesian palm oil value chain, assessing four indicator classes: Deforestation (following EUDR definitions), Land Use Change (LUC) Emissions, Biodiversity (using canopy height heterogeneity as a proxy), and Water Stress (using a meta-indicator combining soil moisture, precipitation, and evapotranspiration anomaly).

The study's findings reveal a significant concentration of environmental risk, with the bottom 5% of facilities linked to over 73% of post-2020 deforestation and 77.9% of total LUC emissions. This disproportionate impact is concentrated among a few corporate actors, primarily in southern Riau and Jambi. Additionally, intensified, estate-dominated supply sheds in Central Kalimantan and northern Sumatra show low biodiversity scores and significant water stress, highlighting a trade-off between productivity and environmental vulnerability.

These results validate the *Codex Planetarius* premise that a targeted approach focusing on the worst performers is the most efficient path toward systemic change in the palm oil industry.

2 Importance of This Work

The core concept of *Codex Planetarius* is to establish a **minimum performance standard** for environmental impacts in the global food and soft commodity trade. It is a framework designed to build consensus on how to measure key impacts – like habitat loss, biodiversity, and emissions – across entire supply chains. The initiative's importance lies in its focus on targeting the **bottom 10-20% of producers,** who are responsible for the majority of environmental harm, rather than just certifying the top performers. By providing this baseline, it aims to create a system that is more efficient and credible than individual governmental actions, ultimately informing international trade agreements and potentially leading to a standardized, universally adopted system.

Epoch is contributing to the materialization of this standard by proposing a methodology that enables the first mile mapping of commodity supply chains, from plot to first point of aggregation, using a state-of-the-art supply shed delineation approach around pre-identified agricultural facilities that aggregate the raw material. This approach establishes a probabilistic, yet robust, relationship between every commodity-producing area and their associated processing facility, and thus ensures a systematic, granular and representative assessment of environmental performance across all commodity producers.

This approach was piloted for the palm oil value chain of Indonesia, and the results presented in this report. The culmination of this work, although still embryonic in light of the scale of soft commodities industries, is to provide a comprehensive view, both granular to the plot, and aggregated at the agricultural facility level, on the state of sustainability of the palm oil industry in Indonesia across 4 indicator classes:

- **Deforestation**, following the EUDR definitions of deforestation assessment.
- **Emissions**, in particular Land Use Change (LUC) Emissions, which are highly variable from producer to producer.
- **Biodiversity**, using canopy height heterogeneity as a proxy for species diversity.
- **Water Stress**, using a combination of soil moisture, precipitation and evapotranspiration anomaly as part of a meta-indicator.

3 Methodology

3.1 Step 1: Identify Aggregation Facilities (Mills)

The initial step involves identifying the first aggregation points in the commodity value chain, typically mills. These facilities establish a direct link to the surrounding landscape, serving as data aggregation anchors for information inferred from nearby commodity growing areas. This includes data on emissions (land-based and non-land-based), deforestation, biodiversity, water stress, and productivity.

The primary challenge lies in scalably and accurately identifying these facilities across commodity-growing regions, often starting with a limited set of reference locations. This is approached as a similarity search, browsing embedding feature spaces for similar vectors. Epoch is currently prototyping various embedding datasets and machine learning algorithms to optimize facility retrieval in commodity landscapes.

For this pilot study, the Trase facility dataset for Indonesian Palm Oil is utilized. This comprehensive dataset provides a list of palm oil facilities in Indonesia, enriched with metadata such as mill name, corporate ownership, processing capacity, operating status, and establishment year. This dataset will be used to demonstrate the *Codex Planetarius* concept for one country and commodity, before outlining a strategy for broader application.

The Trase facilities underwent a quality screening process based on various criteria. The facility confidence score indicates the likelihood of the provided location corresponding to an actual palm mill. Despite Trase being a highly curated dataset, assessing its accuracy is crucial, as precise facility locations are fundamental to this methodology for ensuring the representativeness of derived supply sheds.

3.2 Step 2: Derive Supply Sheds

Once aggregation facilities are identified and characterized with metadata (company, address, commodity), a supply shed is generated using a travel-time based isochrone.

A naive Euclidean distance buffer around facilities does not accurately represent sourcing areas. Instead, a road network-based travel time, with allowances for landscape permeability where roads are absent (considering slope steepness, land cover, and water obstacles), provides a much more accurate approximation of the supply shed. This probabilistic approach, while not 100% accurate, significantly improves upon other methods.

The process is highly scalable, and travel time can be fine-tuned per commodity to determine a realistic sourcing distance, which varies based on topography, road conditions, and freshness requirements.

The first major computation to define supply shed boundaries is the friction layer, which determines the "cost" of movement for each pixel in seconds. The second computation is the cumulative cost calculation, which is the least-cost accumulation of T (friction layer) along any path from the origin to pixel x. This involves evaluating all possible routing options from the origin and assigning the minimum (least-cost) value for each route.

While most parameters generalize well across commodities and jurisdictions, travel time is the most critical parameter for generating representative and accurate supply sheds. For oil palm, a 24-hour freshness window is well-documented to minimize the increase of Free Fatty Acids (FFA), which rapidly occurs after harvesting, especially if the fruit is bruised. High FFA content degrades crude palm oil quality and value. Speed from harvest to processing is essential, and understanding the typical cycle and isolating the transport stage is important. While the entire harvesting cycle typically takes up to 12.5 hours, actual driving time is usually up to 120 minutes (the upper bound of transport), which is the conservative value used to represent the sourcing radius from mill to farm.

Table 1: Summary of the different stages and their corresponding typical duration from tree harvest to processing mill. We want to break it down to isolate the transport component, which is relevant for supply shed delineation.

Stage	Key Activities	Typical Duration	Peer-Reviewed Source
1. Harvesting & In-Field Handling	The FFB is cut, fronds cleared, loose fruits gathered, and the bunch is moved to a collection point.	0.5-1.5 hours	**Saibani et al., 2015**, a time and motion study on FFB harvesting, identified the standard time for all core activities (cutting, collecting, etc.), which fall within this range per tree before any transport waiting begins.
2. Waiting at Collection Point	The FFB sits at the collection point (in-field or roadside) waiting for a truck to arrive for loading.	1-5 hours	This specific waiting period is a component of the overall logistical delay. **Silalertruksa et al., 2012** studied the Thai palm oil industry and noted that transportation from the farm to the mill is a significant source of inefficiency, with this waiting period being a primary, though highly variable, factor.
3. Transport (actual driving time)	The FFB is on a truck, physically moving from the plantation to the processing mill.	**0.5-2 hours**	**Lim et al., 2012**, in a report for the International Food Policy Research Institute, analyzed smallholder supply chains in Indonesia. The report details that farmers may travel up to 50 km to deliver FFB to mills. A journey of this distance, over a mix of poor-quality and paved roads, logically corresponds to a driving time within this range.
4. Queuing & Delay (on road or at mill)	Time the truck spends stationary due to traffic, poor road conditions, or waiting in line at the mill to be weighed and unloaded.	1.5-4 hours	This represents the total non-productive time in transit and at the destination. **Julyanda et al., 2024** identified an "average delay of 3–4 hours" in the overall transport process. Additionally, **Oguoma & Nwafor, 2019** calculated a specific mill queuing time of 1.68 hours using queuing theory.
Total	**N/A**	**4-12.5 hours**	**N/A**

Using travel time aligns with the freshness requirements of raw agricultural produce, as time, not distance, is the determining factor for product quality and economic viability. This travel time can result in a variable distance depending on road network quality and abundance, making the supply shed delineation highly adaptable to local contexts.

The final step involves overlaying the Global Forest Watch's Oil Palm Concessions dataset with the generated supply sheds. If a concession boundary overlaps more than 50% with the delineated supply shed, the concession boundary is used to conflate the supply shed's extent. This ensures optimal representativeness, especially for large estates that might fall outside the typical 120-minute travel time or where incomplete road network data might lead to underestimation of actual travel paths. While this step may not be universally applicable, for oil palm, where estate production dominates, including estate boundaries is highly relevant.

Figure 1 (below) illustrates an example of a supply shed generated using this process for a palm processing mill in West Papua.

Figure 1: Example of a supply shed delineation for a palm oil mill (yellow star in the centre) in West Papua, Indonesia.
The yellow line is the conflated road network from Google Maps API and Overture Maps, and the red outline the concession boundaries as per the Global Forest Watch's Oil Palm Concessions dataset. The concession boundaries are used as an overlay to the delineated supply shed (cyan outline), to ensure full inclusion of concession areas which have more than a 50% overlap with the delineated supply shed.

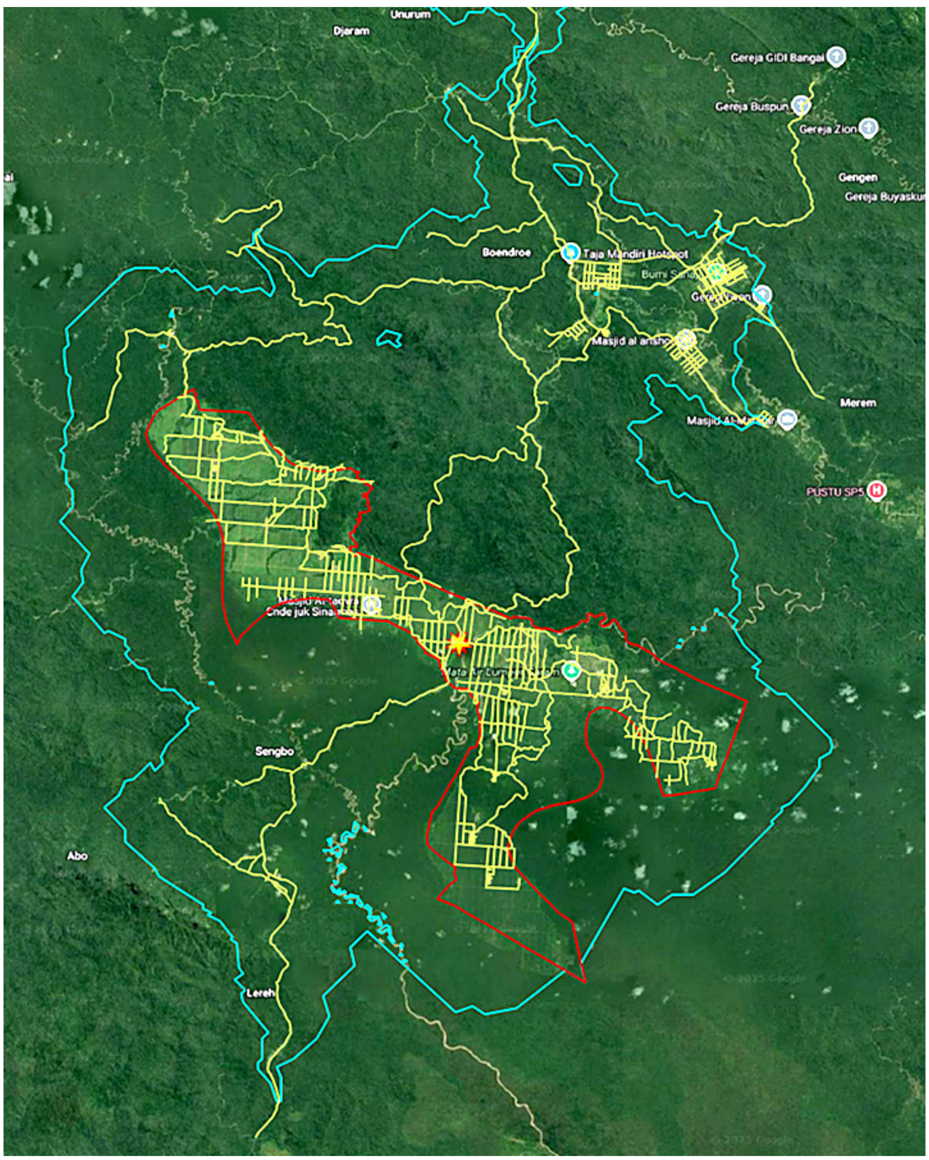

3.3 Step 3: Identify Commodity Production Areas

Commodity maps are a relatively new data type, emerging in response to the EU Deforestation Regulation and the need to identify deforestation-free commodity growing areas. Accurate mapping of commodities, especially perennials like rubber, oil palm, cocoa, and coffee, is crucial for compliance, preventing misidentification of forest cover loss in these systems as deforestation.

Despite their origin in an EUDR regulatory context, these datasets are highly suitable for soft commodity supply chain analytics, enabling exhaustive identification of commodity-producing areas within supply sheds.

The Forest Data Partnership provides initial regional models covering major origins for rubber, oil palm, coffee, and cocoa. For soy, primarily grown in deforestation-exposed Latin America, the GLAD team's Soy Planted Area dataset is annually updated. For cattle, the Global Pasture Watch initiative from WRI provides a global overview of cultivated versus natural pastures, with cultivated grasslands attributed to cattle production (in-situ feeding or livestock feed).

Table 2: Available off-the-shelf datasets to get started with commodity production areas mapping

Commodity	Dataset/Initiative	Coverage	Update Frequency
Rubber, oil palm, coffee, cocoa	Forest Data Partnership preliminary models	Primary origins (regional)	Annually updated
Soy	Soy Planted Area dataset (GLAD team, Song et al, 2021)	Latin America (deforestation-exposed production)	Annually updated
Cattle	Global Pasture Watch (WRI)	Global (cultivated vs. natural pastures)	Annually updated
Timber (plantation forests)	Canopy-covered areas that are not belonging to other commodity layers, and that are not classified as natural forest (Natural Forests of the World 2020)	Presumed plantation forests	Depends on annually updated products

For timber from plantation forests, canopy-covered areas not classified as other commodities or natural forests are presumed to be plantation forests. This currently represents the best available data for identifying timber plantations, and any land clearance (clear-cutting or thinning) of perennial commodity stands near a timber yard is likely attributable to it.

For commodity mapping gaps, Epoch generates its own commodity maps using proprietary labeled datasets specific to commodity and region. Typically, one model per commodity and agro-ecological zone (used for geographic stratification) is trained on AlphaEarth Embeddings data to create an ensemble of commodity models adapted to local agro-climatic contexts.

Whether in-house or off-the-shelf, commodity maps are essentially probability rasters (0-100% probability of commodity presence) that must be vectorized into polygons. To convert the probability field to vectors, specific assumptions are used.

Figure 2: Illustration of the vectorization process of palm probability (green translucid pixels) to a finite vector bound. Notice the effect of the applied parameters (probability threshold, minimum mapping unit, simplification tolerance) on the outcome. Some small palm clusters are left out of the vectorization process because they are small and noisy, and likely do not map to an actual palm stand.

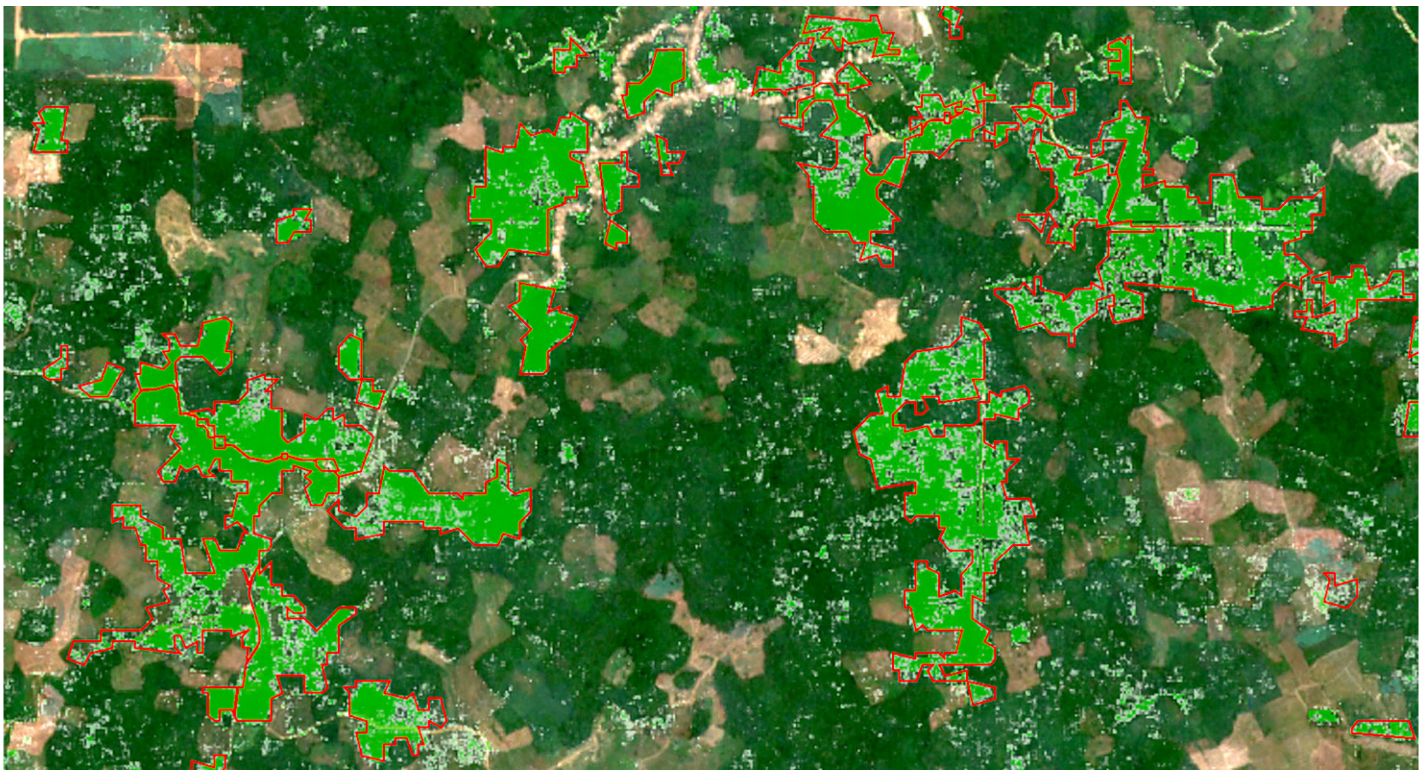

3.4 Step 4: Produce Environmental Metrics

Once individual commodity plots are identified, they are used to compute environmental metrics and associate them with physical aggregation facilities using supply shed delineation. This analysis is grounded in a framework linking supply chain risks to quantitative, satellite-derived indicators.

This matrix serves as a conceptual roadmap, clarifying the rationale behind each chosen metric before detailing its calculation methodology in subsequent sections. All environmental metrics are stored at the plot level, allowing for aggregation using different spatial boundaries (e.g., supply shed, administrative region, country).

3.4.1 Risk-Indicator Framework

A risk-indicator framework was developed to guide environmental metric selection, ensuring robust and transparent analysis. This matrix directly links high-level risks (e.g., monoculture pressure) to measurable proxy indicators (e.g., Canopy Height Heterogeneity), providing the scientific and operational rationale for each metric in this study.

Table 3: Tropical Soft Commodity Risk- Proxy indicator matrix

Risk Category	Specific Risk Factor	Proxy Indicator	Rationale
Climate & Emissions	Land Use Change Emissions	Deforestation Area (Post-2020) & Above-Ground Biomass Loss on converted land	Directly quantifies the primary drivers of LUC emissions, aligning with regulatory standards like the EUDR.
	On-Farm & Processing Emissions	Smallholder vs. Estate Ratio & Presence of Methane Capture Systems (Biogas Domes)	Differentiates between production systems with different farmgate emission profiles and identifies mills that mitigate high-impact processing emissions.
Biodiversity & Ecosystem Health	Habitat Loss & Fragmentation	Landscape Connectivity Score (e.g., from Omniscape)	Measures how fragmented the natural landscape is, which is a direct proxy for habitat quality and the ability of species to move and thrive.
	Monoculture Pressure & Pest/Disease Vulnerability	Canopy Height Heterogeneity (Rao's Q Index)	A more uniform canopy indicates a monoculture, which is more susceptible to pest and disease outbreaks like bagworms and Ganoderma.
Water & Hydrology	Drought & Water Scarcity	Meteorological Water Stress Index (combining soil moisture, precipitation/PET ratio, and evapotranspiration anomaly)	Captures immediate climate-driven water shortages that can reduce yields and stress plantations.
	Operational Disruption (Floods/Drought	Hydrological Flow Dynamics (Baseflow & Quickflow from a model like SWAT+)	Quantifies the landscape's response to rainfall, directly assessing the risk of flash floods that sever supply routes or river droughts that halt mill operations.
Land & Soil Health	Long-Term Productivity Loss	Soil Erosion & Sediment Yield (from a model like SWAT+)	Measures the degradation of the core productive asset (topsoil), which impacts long-term yields and requires costly interventions.

3.4.2 Deforestation

Deforestation assessment quantifies forest cover loss within specified areas and monitoring periods, adhering to the European Union Deforestation Regulation (EUDR) definition for regulatory compliance and standardized reporting.

The methodology aligns with the EUDR definition of deforestation as "the conversion of forest to agricultural use, whether human-induced or not," encompassing:

1. EUDR Forest Definition (Table 4, below).

Table 4: EUDR Forest Definition Criteria

Criteria	Requirement
Tree canopy cover	Greater than 10%
Tree height	Minimum 5 meters at maturity
Minimum area	0.5 hectares
Exclusions	Areas under agricultural or urban land use

1. Deforestation Criteria (Table 5, below).

Table 5: EUDR Deforestation Criteria

Type	Description
Land Use Conversion	Conversion of forest land to agricultural use (cropland, pasture, plantation agriculture)
Canopy Removal	Permanent removal of forest canopy through human activities
Ecosystem Loss	Loss of forest ecosystem functions and services
Temporal Threshold	Deforestation events occurring after December 31, 2020 (EUDR cut-off date)
Exclusions	Natural forest disturbances (wildfires, storms, disease) that don't result in permanent land use change to agriculture

Under EUDR, natural forest disturbances (wildfires, storms, disease) not resulting in permanent land use change to agriculture are not classified as deforestation. The goal is to isolate deforestation attributable to commodity production, excluding external factors, to assess producer environmental performance.

Primary change detection relies on a combination of advanced algorithms and datasets, including Epoch's Dynamic World-CCDC Change Detection, RADD (Radar for Detecting Deforestation), GFW-GLAD Annual Forest Loss, GLAD-S2 (Sentinel-2 based alerts), DIST (Disturbance Alert), JRC-TMF Annual Deforestation, and Epoch's Biomass Change Trends.

The system integrates these multiple alert sources using a ruleset that combines various alerts with confidence levels. This provides comprehensive forest loss detection through agreement analysis between methods.

The natural Forest Reference Masks against which forest change detection is compared are the JRC-TMF2020 Forest Mask, Google DeepMind Natural Forests of the World 2020, and Forest Data Partnership Commodity Probability Models.

Forest loss pixels are identified where annual loss events occur within the specified monitoring window and intersect natural forest areas meeting EUDR forest criteria. Deforestation area calculations use weighted pixel proportions to avoid over- and under-accounting. Agreement analysis between multiple detection methods provides confidence metrics, with agreement between Dynamic World-CCDC and GLAD detection representing at least 25% of detected non-compliance area for high confidence assessments.

Five confidence levels are generated based on multiple criteria. These confidence levels are further categorized into "critical" (medium to very high), "non-critical" (low and very low), and "no deforestation" (no non-compliance area detected for the given plot).

These confidence levels are further categorized into “critical” (medium to very high), “non-critical” (low and very low) and “no deforestation” (non compliance area detected is 0 for the given plot).

Table 6: Deforestation Confidence Levels

Level	Criteria	Category
Very High	Non-compliance area >10% of plot area OR >1 ha (whichever smaller), agreement between datasets ≥25% of detected area, plot overlaps with natural forest or protected areas	Critical
High	Non-compliance area >10% of plot area OR >1 ha (whichever smaller), plot overlaps with natural forest or protected areas	Critical
Medium	Non-compliance area >10% of plot area OR >1 ha (whichever smaller)	Critical
Low	Non-compliance area >5% of plot area OR >0.5 ha (whichever smaller) OR has forest overlap	Non-critical
Very Low	Non-compliance area > 0	Non-critical
No Deforestation	Non-compliance area = 0	No deforestation

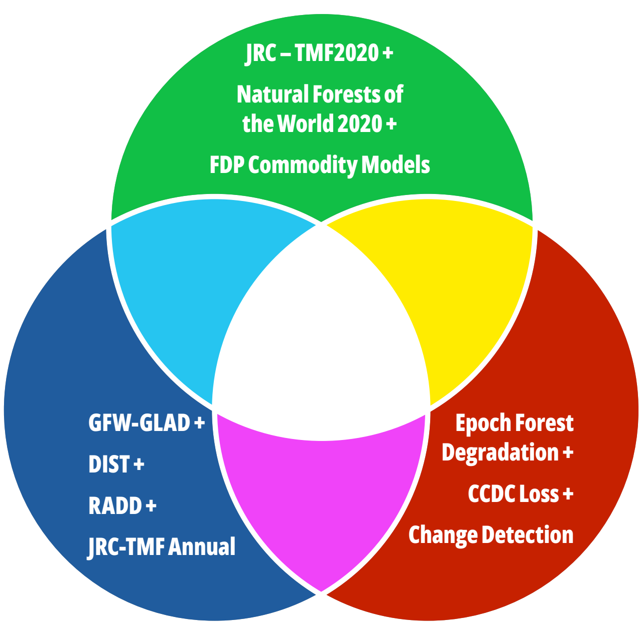

Figure 3: Convergence of evidence approach whereby the off-the-shelf detection products (blue) and the Epoch’s change detection approach (red) is overlaid with the natural forest mask (green) to identify deforestation taking place over natural forest. While cyan and yellow mean that either data source identified deforestation, white identifies convergence in the detection from both data sources.

3.4.3 LUC Emissions

Land-use change (LUC) emissions assessment quantifies CO2-equivalent emissions from above-ground biomass loss using spatio-temporal biomass modeling. This methodology differentiates between LUC emissions (from land cover changes) and non-LUC emissions (from stable land cover with agricultural practices), providing comprehensive carbon accounting for supply chain sustainability assessment.

The primary data sources are GEDI LiDAR L4A Footprints and AlphaEarth Satellite Embeddings.

Biomass density is modeled annually from 2017 onward using machine learning approaches trained on GEDI reference data. Above-ground biomass (AGB) estimates are converted to CO2-equivalent emissions using standard biomass-to-carbon conversion factors and the molecular weight ratio.

A land cover transition analysis is overlaid with emissions data layers to categorize biomass-related emissions into the following two categories: AGB+BGB Emissions (pixels undergoing land cover change) and Biogenic Emissions (pixels with stable land cover). SOC Emissions are characterized using emissions factors resulting from land use conversion from forest to palm plantation.

In a future iteration, to align with the Greenhouse Gas Protocol (GHGP) Land Sector Removal Guidance (LSRG) emissions accounting principles, emissions estimates will be produced from the year 2000 onwards. This will ensure that emissions identified up to 20 years ago are accounted for appropriately using the linear discounting principle.

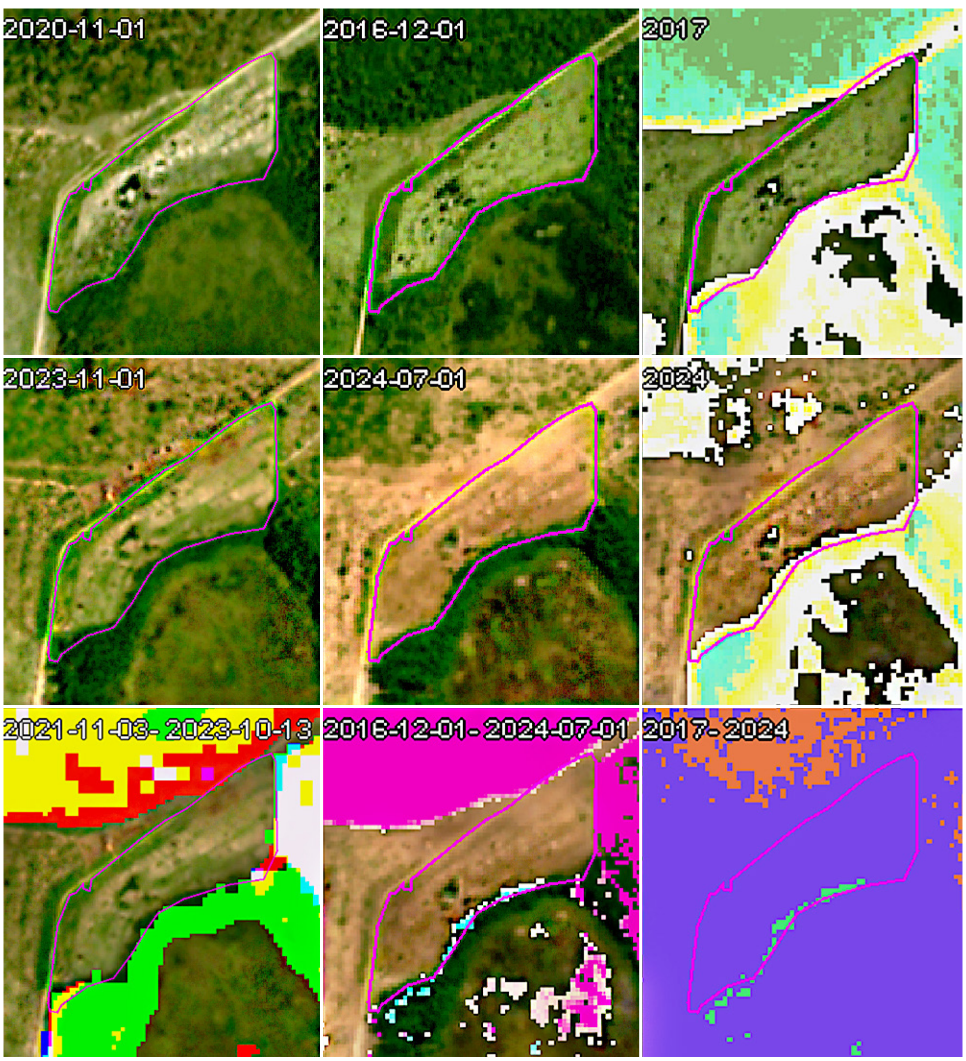

Figure 4: Illustration of deforestation and LUC emissions for a given plot. Left: the plot in question would not be considered to be deforested, because the deforestation occurrence is marginal (i.e. a few pixels on the edge of the plot border, which corresponds to < 5% of plot area). Center: decrease (red pixels) and increase in biomass stock (blue pixels). Right: the categorization of emissions according to a land cover transition mask where purple is biogenic emissions (no land cover change in the monitoring period), red is positive LUC emissions (biomass loss) and green is negative LUC emissions (biomass gain).

3.4.4 Non-LUC Emissions

For non-LUC emission sources, this approach relies on farmgate emissions factors for smallholder and estate palm production. These factors represent emissions from on-farm practices and inputs used in plantations.

The distinction between smallholder and estate is made using specific criteria. It is important to note that a "plot" is considered a continuous homogeneous area growing palm, not a delineation based on ownership, as remote sensing data cannot make that distinction.

The following emissions factors were used to characterize non-LUC emissions in oil palm plantations:

Table 7: Non-LUC Emissions Factors for Oil Palm Plantations

Region	Estate Factor (tCO2e/ha/year)	Smallholder Factor (tCO2e/ha/year)	Activity	Source(s)
Indonesia	2	1.4	Fertilization (N2O, CH4, empty fruit bunches, etc)	Taken Safitri et al., 2024

More granular farmgate emissions factors representing different parts of Indonesia are not available in peer-reviewed literature, making it difficult to assess the variability of non-LUC emissions at this stage. Most available information concerns LUC (described in the previous section and incorporating all ecosystem carbon pool changes), and for other farmgate activities, the bulk of emissions comes from fertilization, with other contributions being negligible.

3.4.5 Biodiversity

Biodiversity assessment uses canopy height heterogeneity as a proxy for species diversity, employing Rao's Quadratic Entropy (Rao's Q) index. This approach acknowledges the correlation between structural complexity and species richness in forest ecosystems, providing quantitative biodiversity metrics for supply chain sustainability assessment.

Primary Data Sources used to determine Rao's Q are GEDI LiDAR L2A Footprints and AlphaEarth Satellite Embeddings.

The Rao's Q diversity index is calculated using spatial variance and dispersion of canopy height metrics within a search window. Values below 0.3 indicate low diversity (typically monocultures and plantations), while values above 0.5 to 1.0 characterize fragmented, biodiverse landscapes.

Canopy height heterogeneity is computed from the spatial variance of canopy height model (CHM) derivatives within assessment areas. Annual canopy height means and standard deviations are calculated, with maximum values representing the theoretical potential for mature natural forest. The Rao's Q diversity score integrates spatial heterogeneity patterns as an indicator of landscape diversity.

In the future, biodiversity will be further characterized using a landscape connectivity model that would provide a connectivity score for each supply shed, offering a proxy for natural habitat disruption.

3.4.6 Water Stress

Water assessment provides a comprehensive hydrological analysis through four key indicators, culminating in a composite Water Stress Index. This methodology quantifies water availability, stress conditions, and hydrological anomalies for agricultural monitoring and water resource management.

The primary data sources used to calculate the water stress indicators are CHIRPS Precipitation, MODIS Evapotranspiration, and SMAP Soil Moisture.

The specific sub-indicators calculated are:

1. **Precipitation/PET Ratio:** Measures water input balance against atmospheric demand
2. **Soil Moisture Percentile:** Evaluates current conditions relative to a historical climatology
3. **Evapotranspiration Anomaly:** Measures vegetation water use deviation from historical norms

A 5-year climatology period ending 1 year before the current monitoring period is used. All components are normalized to a specific range with weights summing to 100%. Water bodies and built-up areas are automatically excluded to focus on agricultural and natural landscapes.

From the above three sub-indicators, a Water Stress Index composite indicator is produced.

3.4.7 Overall Risk Score

The Overall Risk Score is a composite indicator designed to assess environmental and sustainability risks across palm oil supply sheds. This methodology combines indicators from previous sections (total emissions, biodiversity, water stress, non-compliance area, commodity/supply shed area ratio, estate/smallholder ratio) into a single normalized score ranging from 0.0 (low risk) to 1.0 (high risk), enabling comparative risk assessment across facilities and supply sheds.

All components use percentile ranking to normalize values to a 0-1 scale, ensuring equal contribution, robustness, and comparability. Equal weighting assumes all components are equally important. While percentile ranking may mask absolute improvements if all facilities improve simultaneously, this approach is fit-for-purpose for ranking them at a given instant in time.

4 Results

The results section is structured in the following way for each indicator:

- Distribution of supply sheds and their indicators on a map and a histogram.
- Bar chart showing the histogram aggregate (sum or mean, depending indicator) per percentile range (0-5th, 5-25th, 25-50th, 50-75th, 75-95th, 95-100th) and their corresponding production capacity (data from the Trase oil palm facilities[1] dataset).
- Pie charts showing the histogram aggregate per percentile range and per company.

Figure 5: Facilities screened for Indonesia and their respective score ("high": green, "medium": yellow). No "low" confidence has been identified for the Trase dataset used.

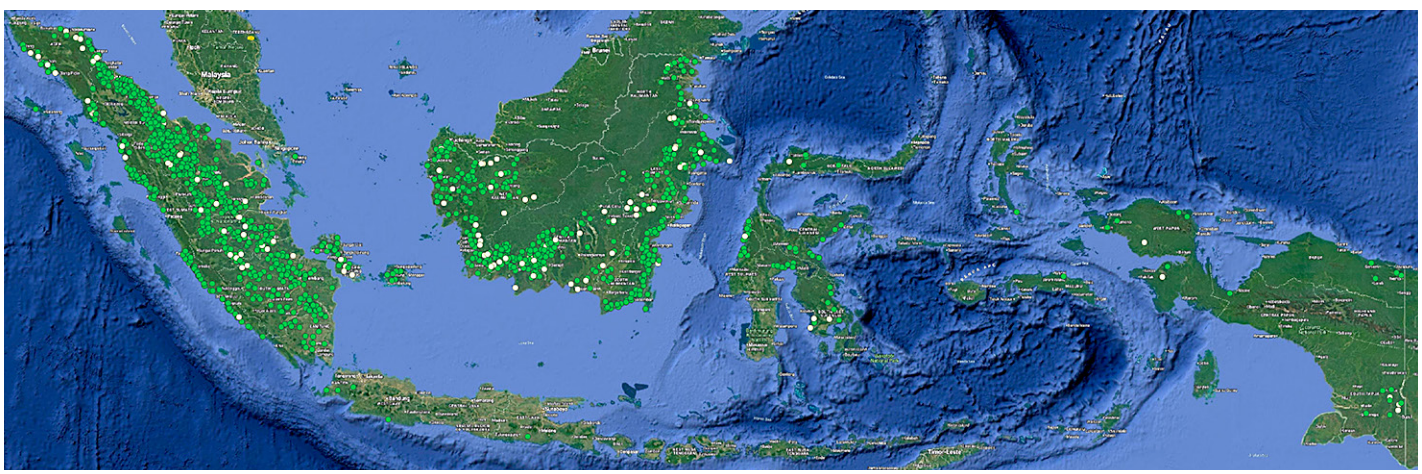

4.1 Facilities Validation

The facilities' location screening classifies the overwhelming majority as "high" confidence (above 0.7 score), with only 67 identified with a "medium" score, and none as "low" score. This is unsurprising given the high quality and curation that went into the production of the Trase dataset. The "medium" score stems from the lack of road and building data in the Overture datasets for rural areas, which inherently brings down the score for certain areas. For future supplier datasets that have not been curated and vetted by a trusted authority, this facility's pre-screening feature may be instrumental in determining whether the datasets are sufficiently accurate to carry out a supply shed analysis.

[1] https://trase.earth/explore/facilities-data/map?facilityTypeId=indonesia-palm-oil-mills

4.2 Supply Shed Characteristics

4.2.1 Supply Shed Statistics

Table 8: Supply Shed and Commodity Growing Area Statistics

Total facilities processed	1,290
Total plots (contiguous homogenous palm growing areas)	859,811
Total commodity growing area (ha)	10,499,973
Of which estate (ha)	8,294,978 (79%)
Of which smallholder (ha)	2,204,994 (21%)
Total supply shed area (ha)	**51,139,526**

The summary data for the supply sheds and the commodity growing area mapping is summarized in Table 8 (above).

Descals et al., 2024 produced similar data for the 1990-2021 period, and came up with 10,072,613 ha and 3,023,232 ha for estate and smallholder areas respectively, and a breakdown of 70/30% for estate/smallholder. The numbers are 20% larger for estate and 30% larger for smallholders, a difference that could be explained by the fact that the epoch of 31 years used to extract information in Descals et al., 2024 could encompass retired plantation areas, whereas this study used a shorter epoch (2017-2024) that takes a shorter "snapshot" of the active oil palm land footprint.

The supply sheds for Kalimantan and Sulawesi, and Sumatra, can be seen in Figure 6 and 7.

Figure 6: Supply Sheds produced for Kalimantan and Sulawesi

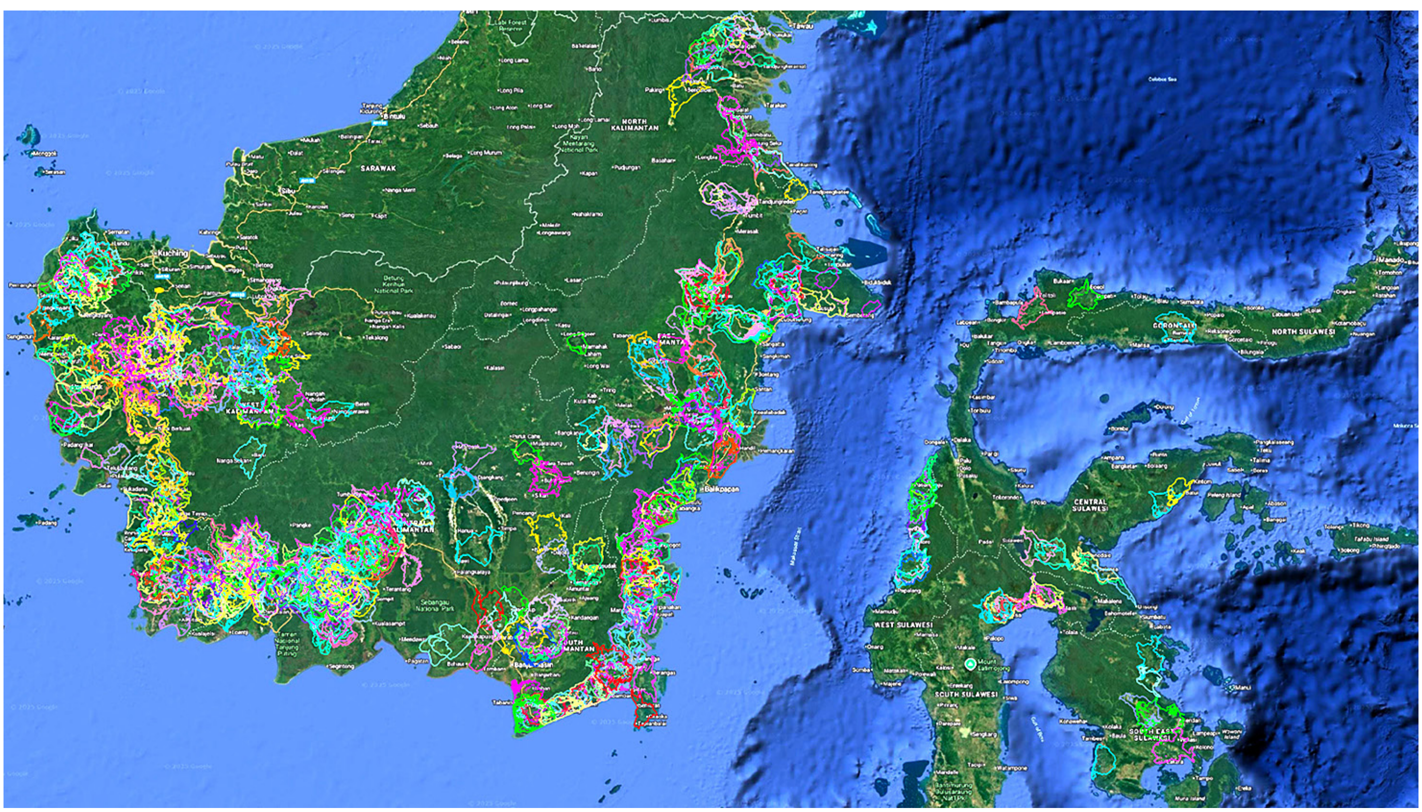

Figure 7: Supply Sheds produced for Sumatra

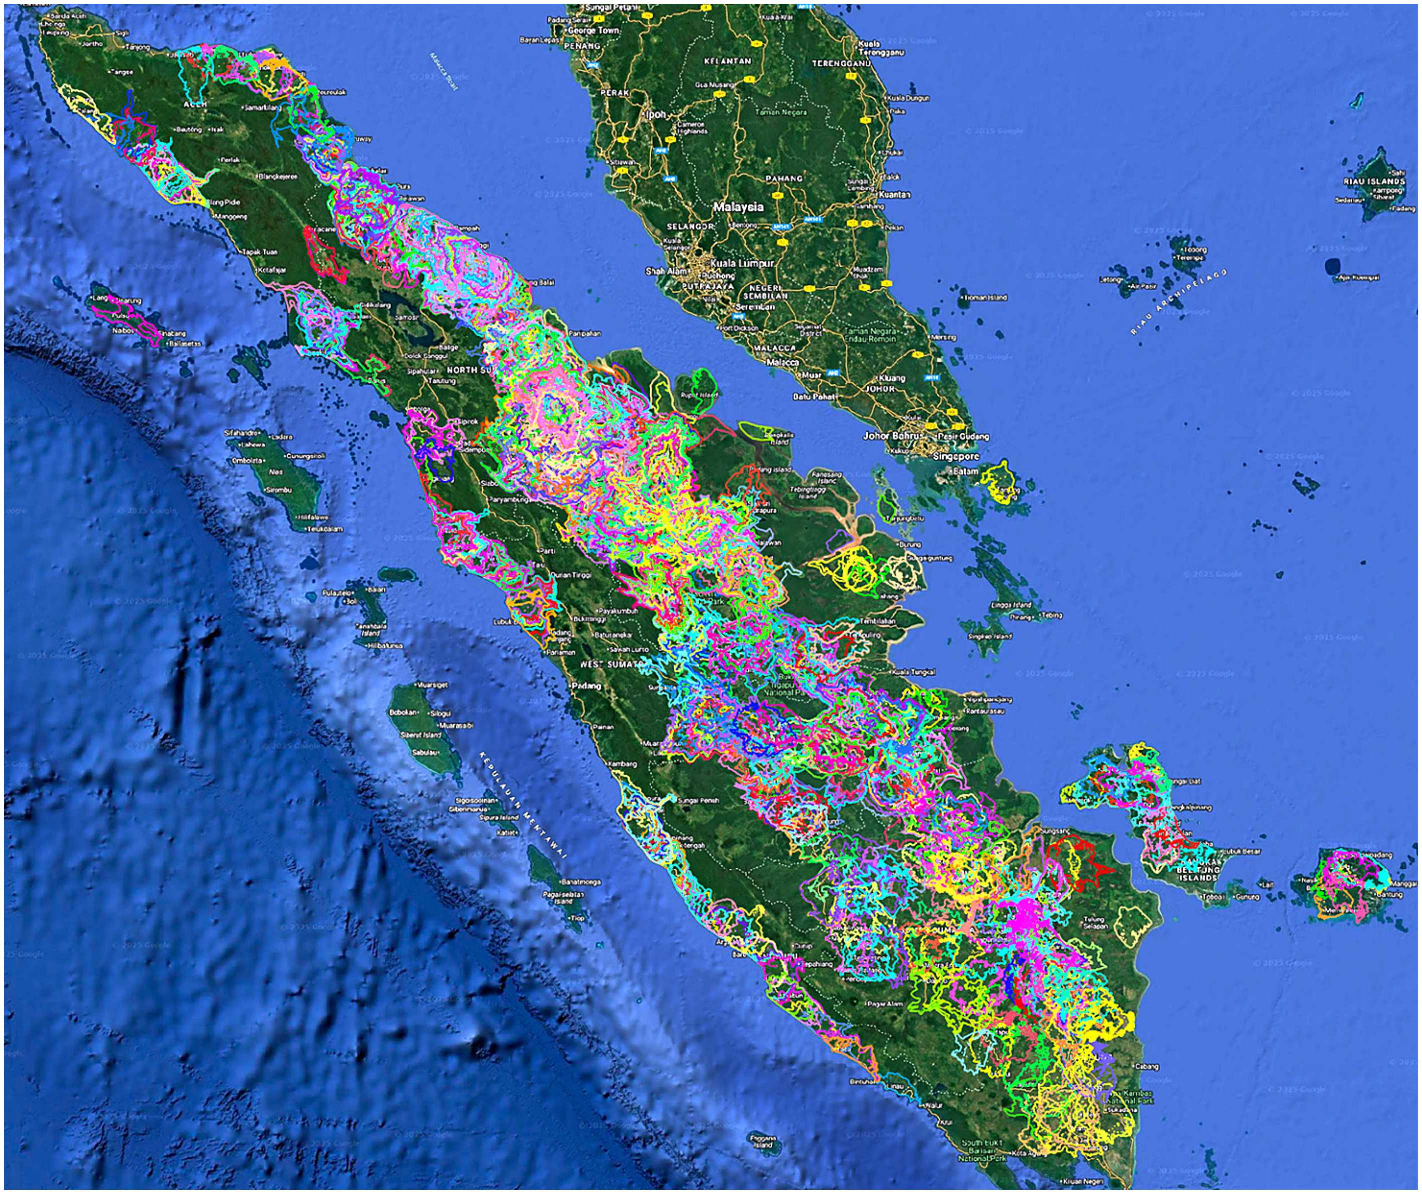

4.2.2 Supply Shed Size

Mapping supply shed sizes on a map reveals that Sumatra has larger supply sheds around the facilities than Kalimantan and other regions. The main reasons for this is the flat topography and the extensiveness and quality of the road infrastructure, allowing the expansion of palm plantations in the plains. Kalimantan is a more recent palm origin, and its development may have been partly constrained by the high prevalence of peat soils, but also because of lower road network density.

Figure 8: Supply shed size (in ha)

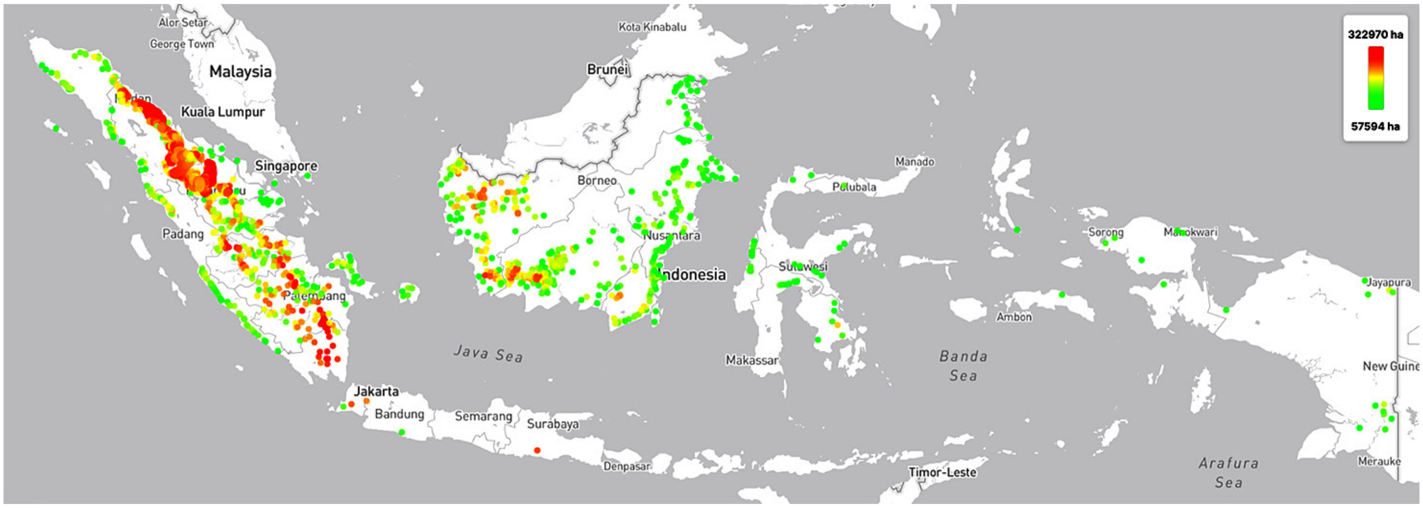

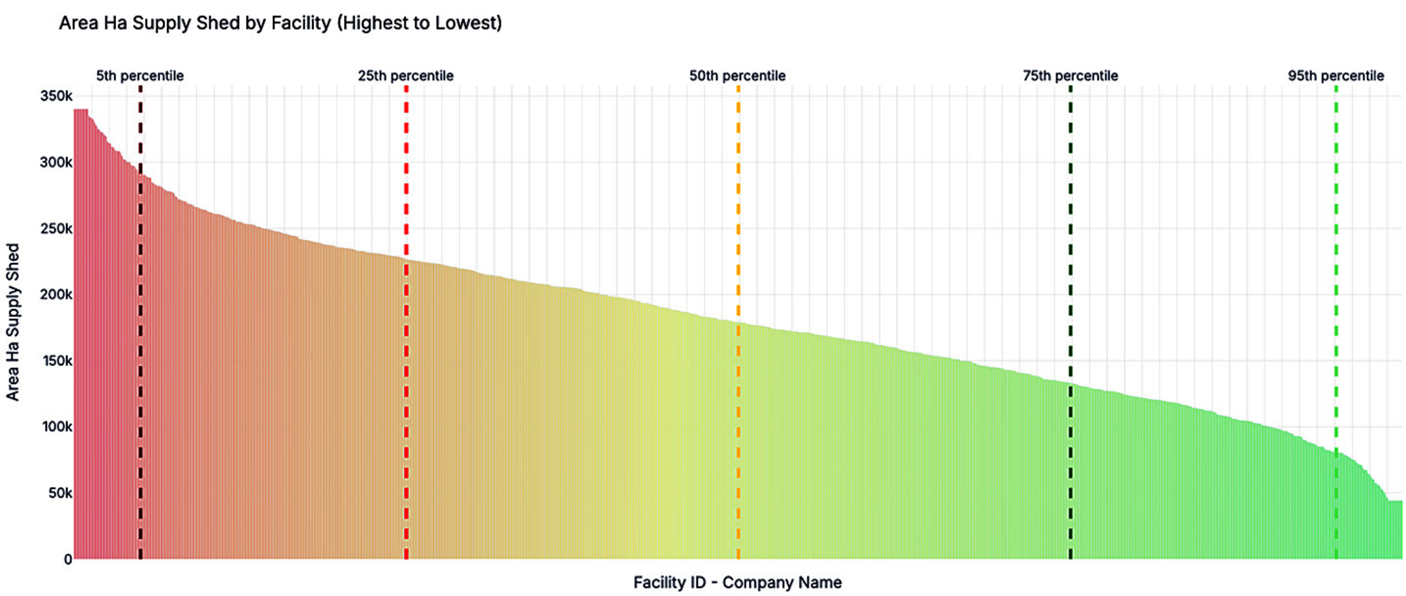

4.2.3 Commodity / Supply Shed Area Ratio

When examining the commodity growing area as a percentage of the supply shed area, two clusters emerge in Northern Sumatra and Central Kalimantan, representing the densest commodity-producing areas in Indonesia. On the western Sumatran coast and western Sulawesi coast, we see small red clusters, which can be explained by the steep topography around those areas, forcing the production of palm on the milder slopes in valley bottoms.

Figure 9: Palm Oil-growing area per supply shed (in ha)

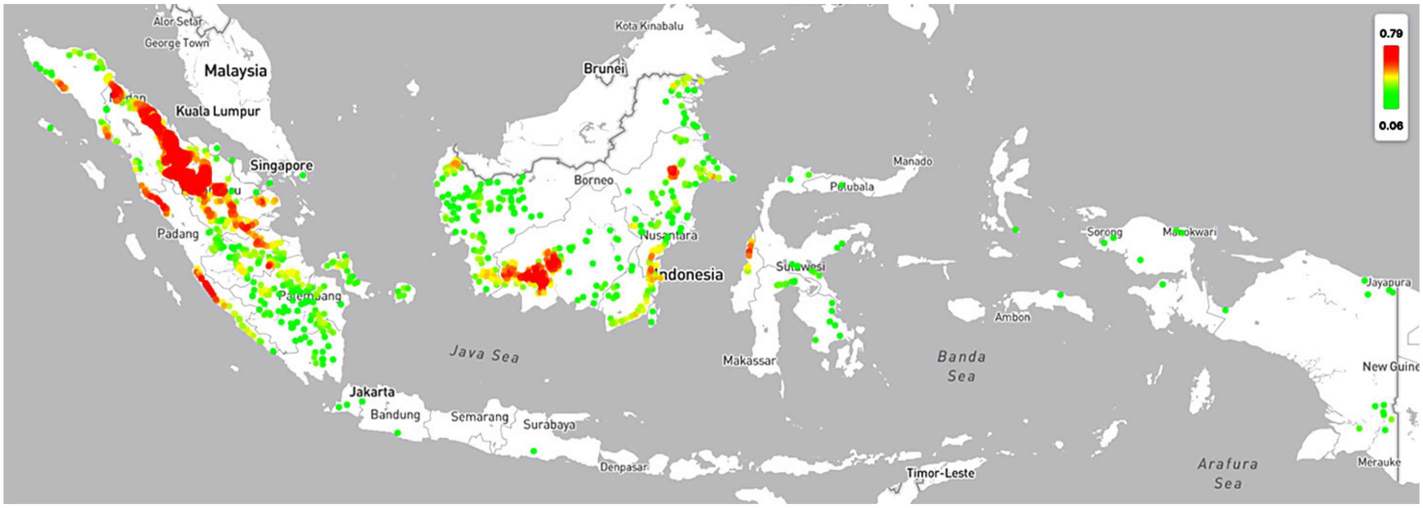

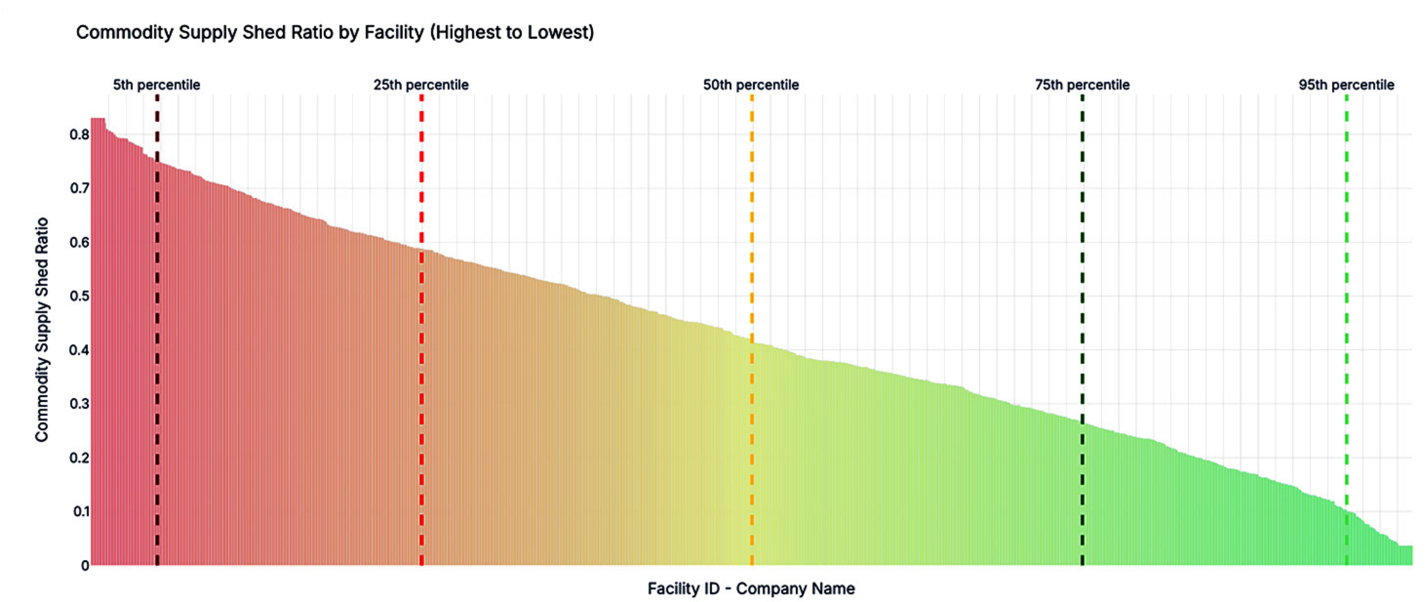

Figure 10: Example of a supply shed on the Western Sumatran coastline, where palm oil is squeezed on the coastal plain, and slowly on the steeper slopes as the supply shed moves inland

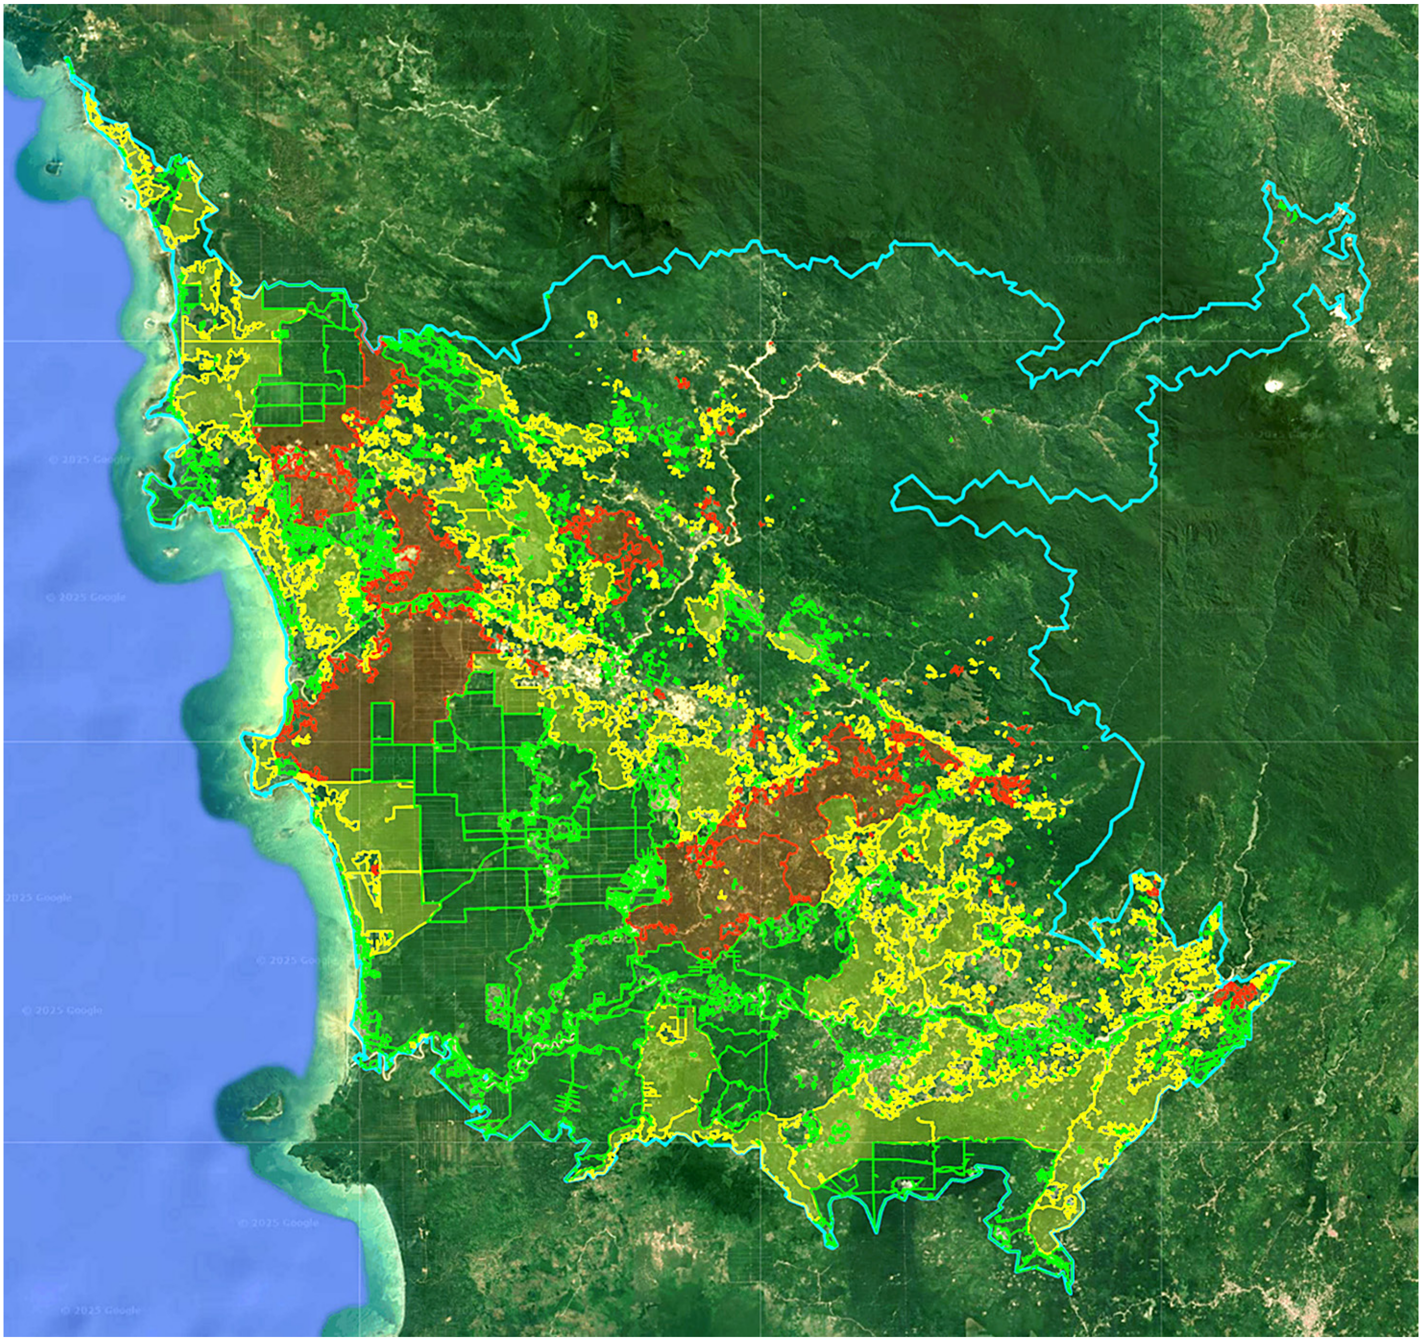

4.2.4 Estate / Smallholder Area Ratio

Across all facilities, the breakdown of estate to smallholder growing area is 79% estate and 21% smallholder, but with significant regional variations. Some areas in northern Riau and Central Kalimantan seem to be dominated by estate structures, with very little smallholder developments. The cluster in Central Kalimantan seems to agree with findings from Potter, 2016, but less so with the smaller cluster in northern Riau. In spite of that, the specific area in Riau identified (north-west of Pekanbaru) is a highly intensive palm production sector, based on the data evidence Epoch has produced for this report.

Typically, the smallholder areas are usually spread around the outskirts of the estates, in areas of variable sizes (but typically smaller than 100 ha). Figure 12 illustrates an example of the locations Epoch has identified as smallholder palm for a given supply shed in East Kalimantan.

Figure 11: Estate / smallholder ratio, where red areas are cases of extreme prevalence of estate to smallholder

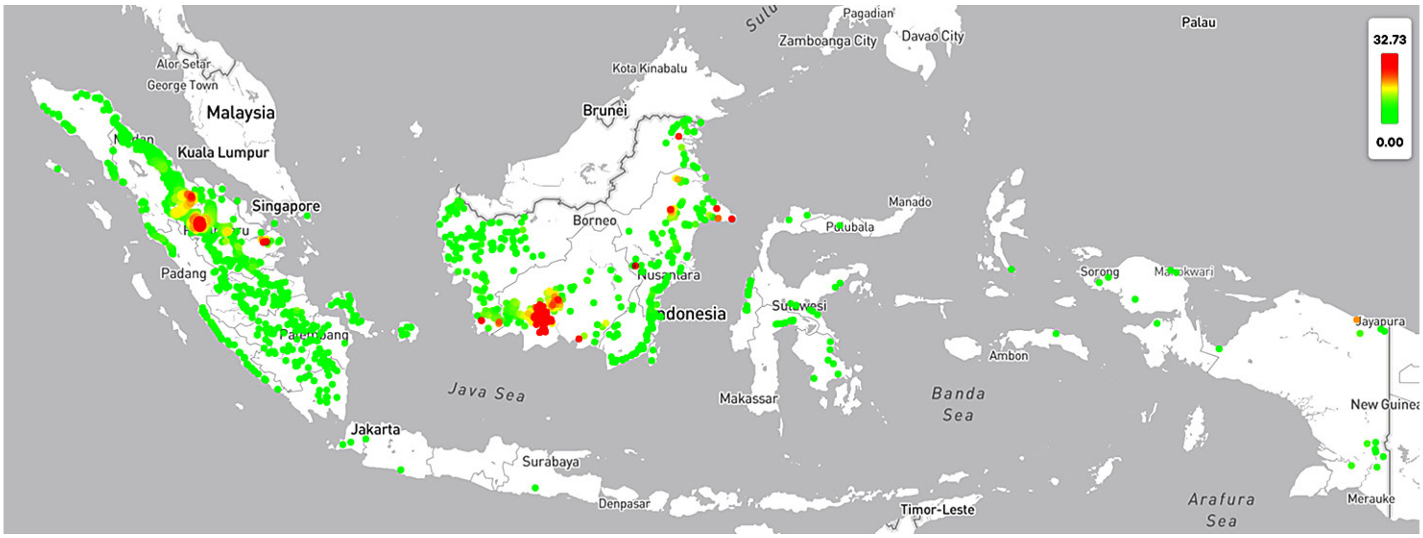

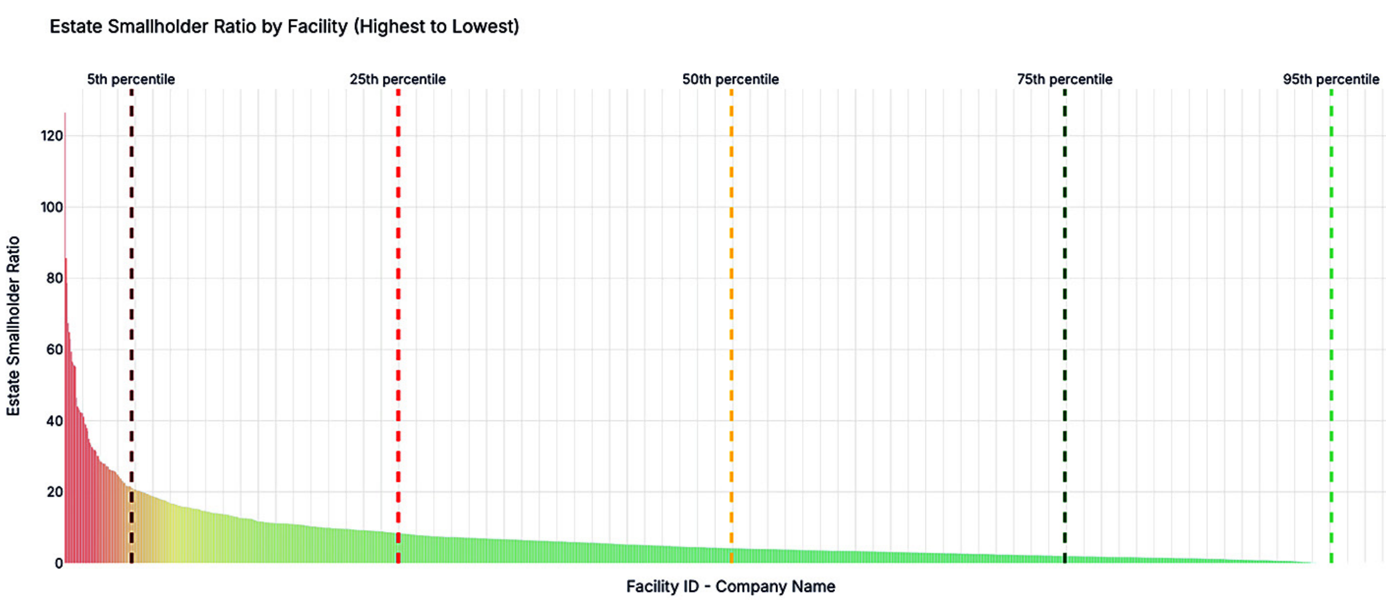

Figure 12: Example of the typical estate (turquoise) and smallholder (khaki) distribution in a supply shed

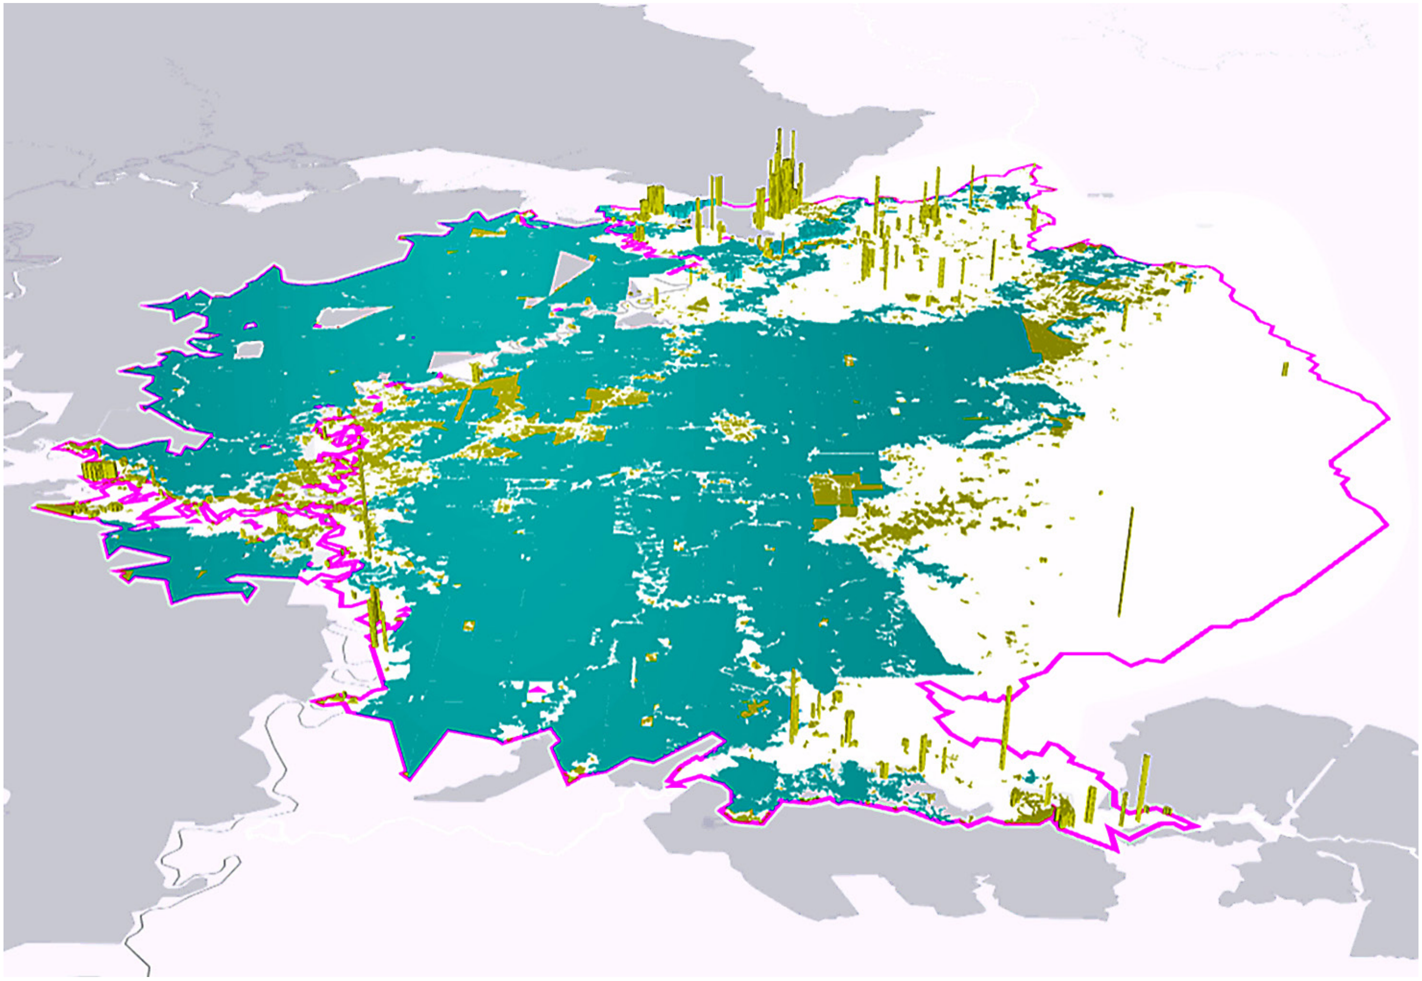

4.3 Deforestation

Deforestation, following the EUDR definition, is concentrated in a very small proportion of the supply sheds, namely in southern Riau and Jambi. In general, it is clear from the palm facilities distribution in Figure 13 that the bottom 5th percentile is responsible for the bulk of these deforestation cases.

Figure 13: Per-supply shed non-compliant deforested area (ha) plotted on the map of Indonesia (top) and on a histogram ranking highest (red) to lowest (green) deforestation area

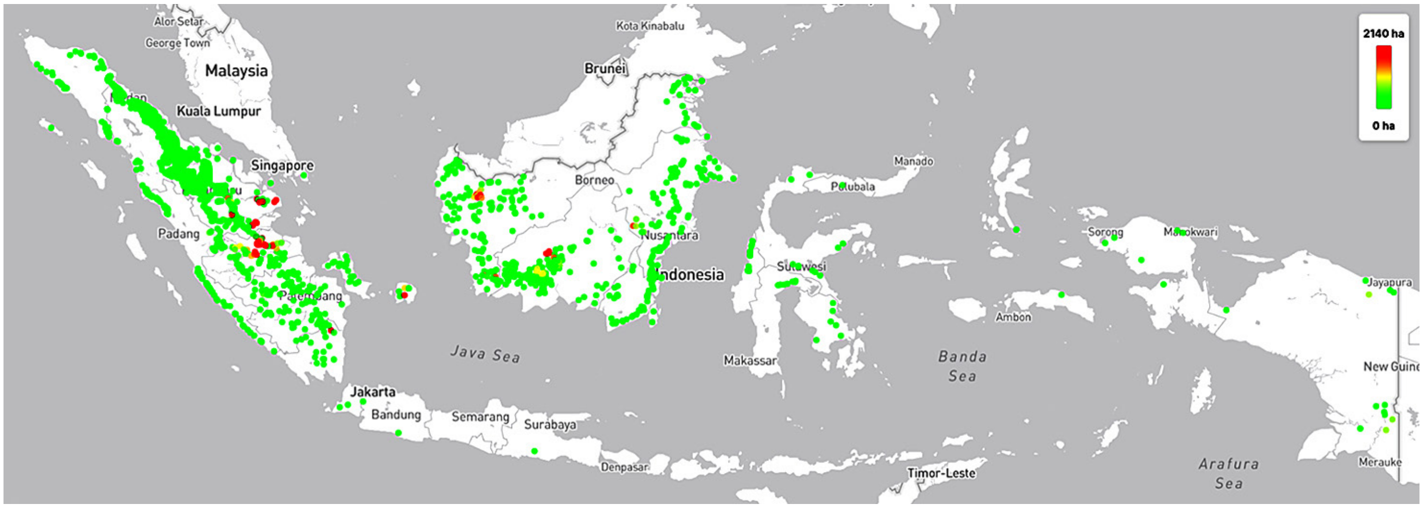

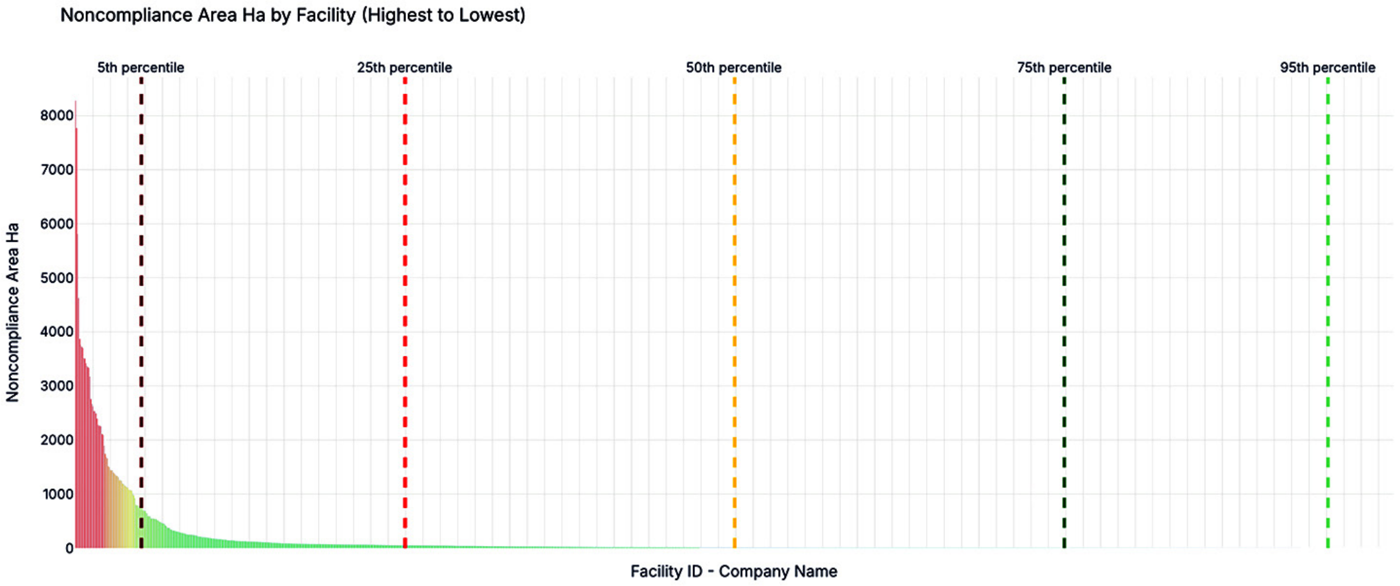

Figure 14: Relative contribution of deforested area per percentile range

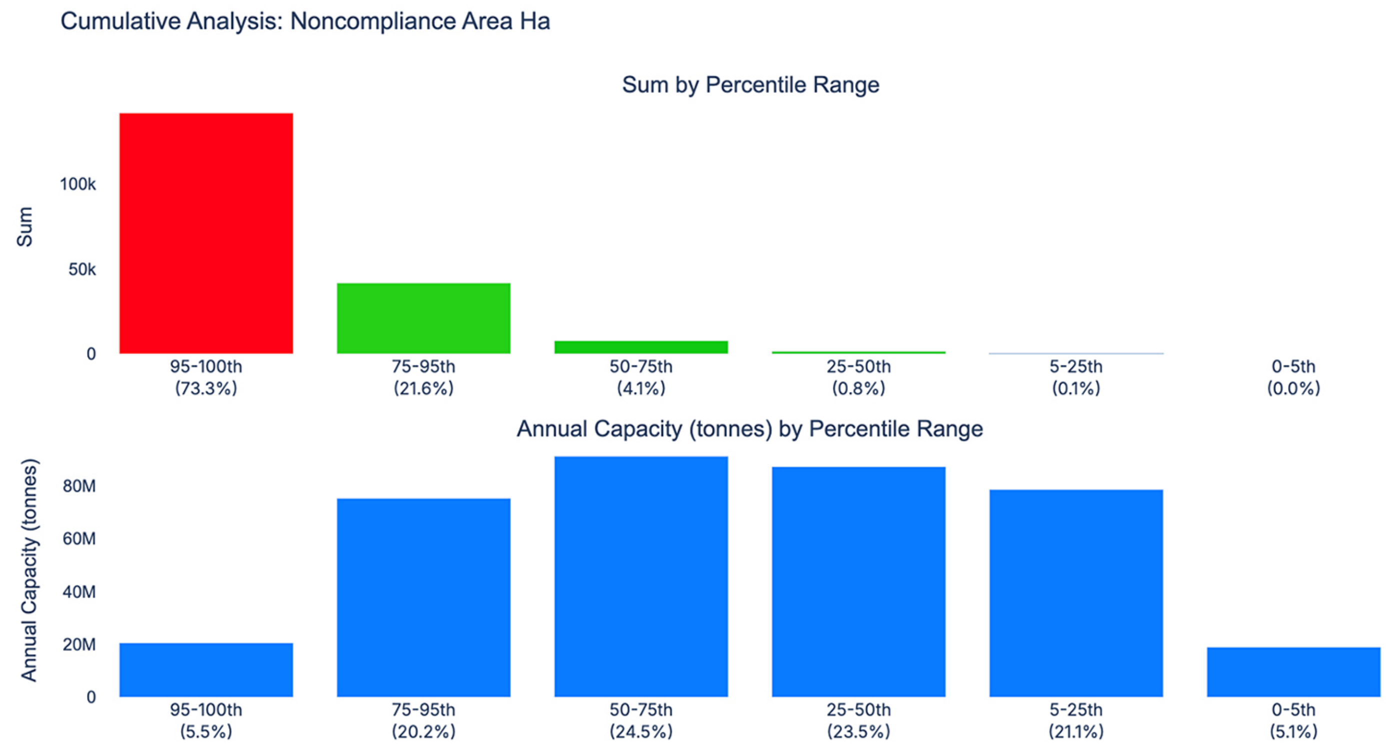

Figure 14 identifies that his bottom 5th percentile contributes to 73.3% of all deforestation, corresponding to 141,436 ha deforested since the EUDR cut-off date. Moreover, this affects a corresponding 5.5% of total annual production capacity (or 20,580,000 tonnes) across all facilities.

Figure 15: Relative contribution of deforested area per percentile range and per company. The size of the hole in the pie chart determines the relative importance of that pie chart for deforestation

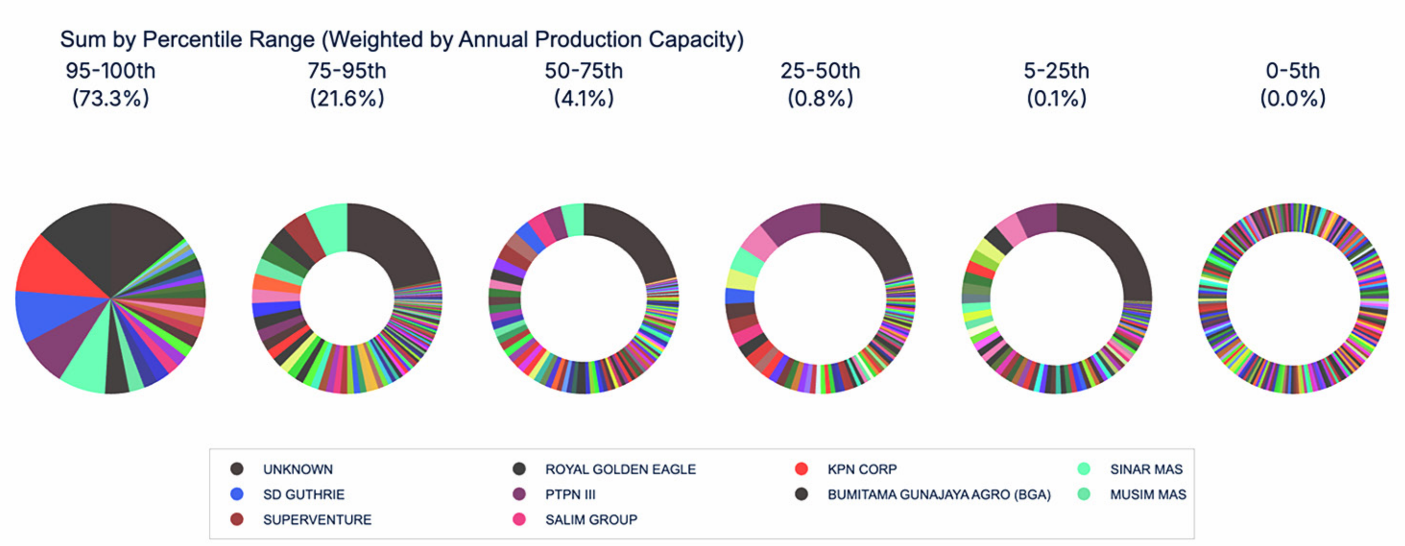

Zooming into this lower 5th percentile again, Figure 15 illustrates that 5 companies (Royal Golden Eagle, KPN Corp, SD Guthrie, PTPN III, Sinar Mas) have exposure to approximately 49% of deforestation (69,402 ha) in that percentile range.

4.4 Emissions

4.4.1 LUC Emissions

LUC emissions are highly correlated with deforestation data, since most of LUC emissions result from the conversion of natural land (most often natural forest) to palm plantations.

Figure 16: Per-supply shed LUC Emissions (tCO2e/year)

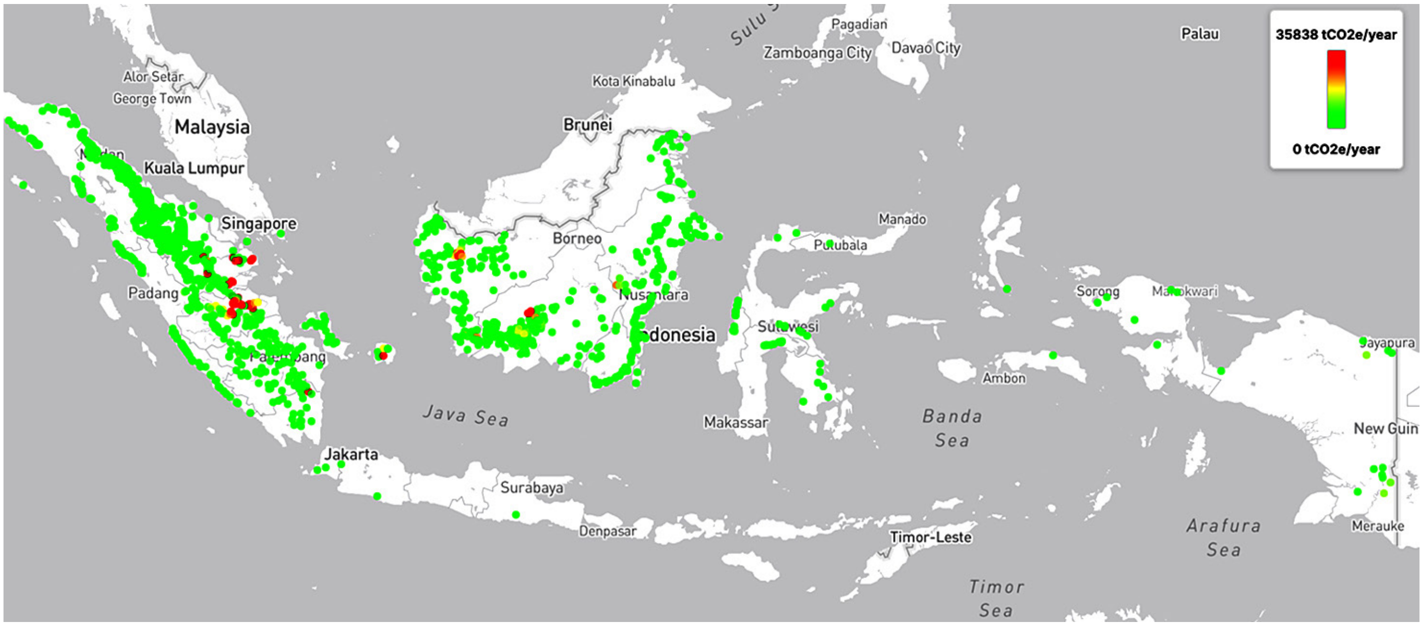

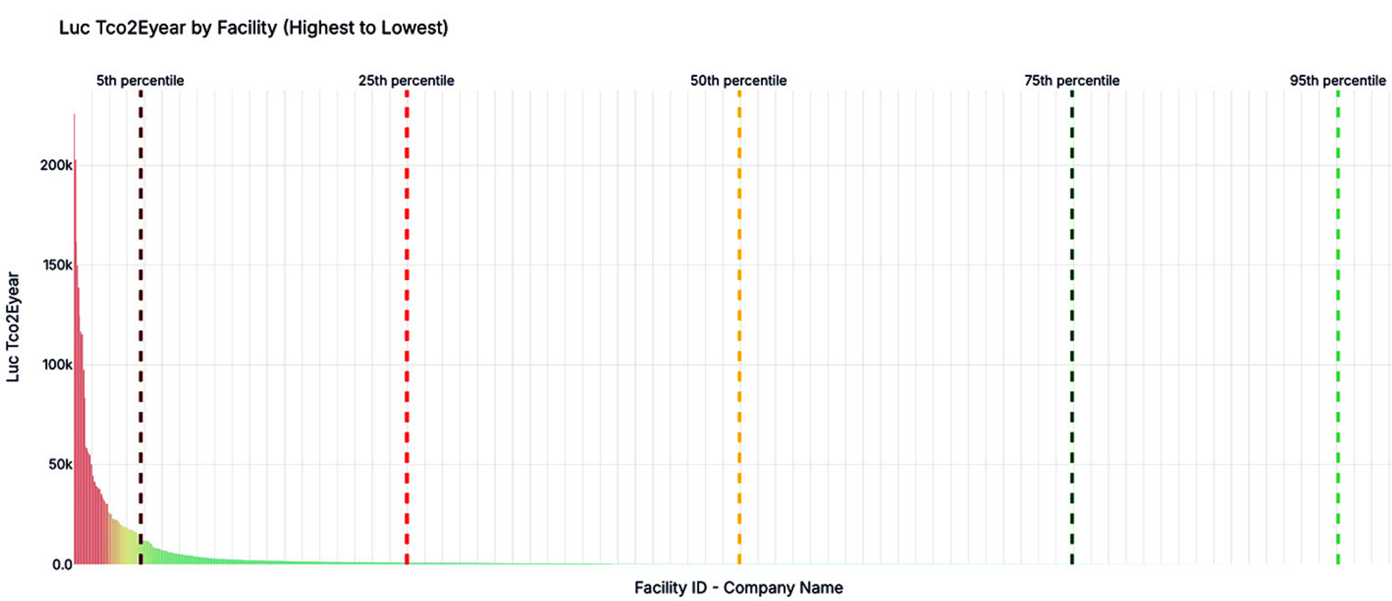

Figure 17: Relative contribution of LUC emissions per percentile range

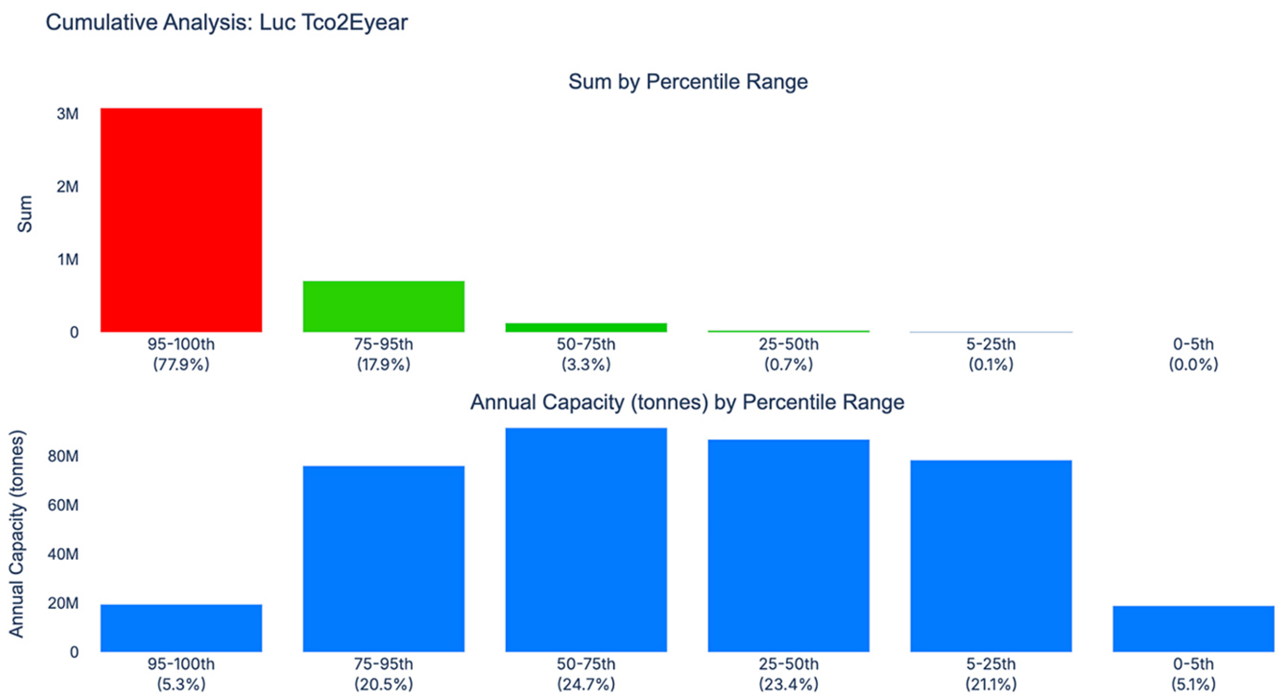

Figure 17 shows that the LUC emissions are even more skewed, with the bottom 5th percentile representing 77.9% of total LUC emissions (3,078,634 tCO2e/year), and corresponding to 5.3% of the total annual production capacity (19,560,000 tonnes/year) across all facilities.

Figure 18: Relative contribution of LUC emissions per percentile range and per company. The size of the hole in the pie chart determines the relative importance of that pie chart for LUC emissions

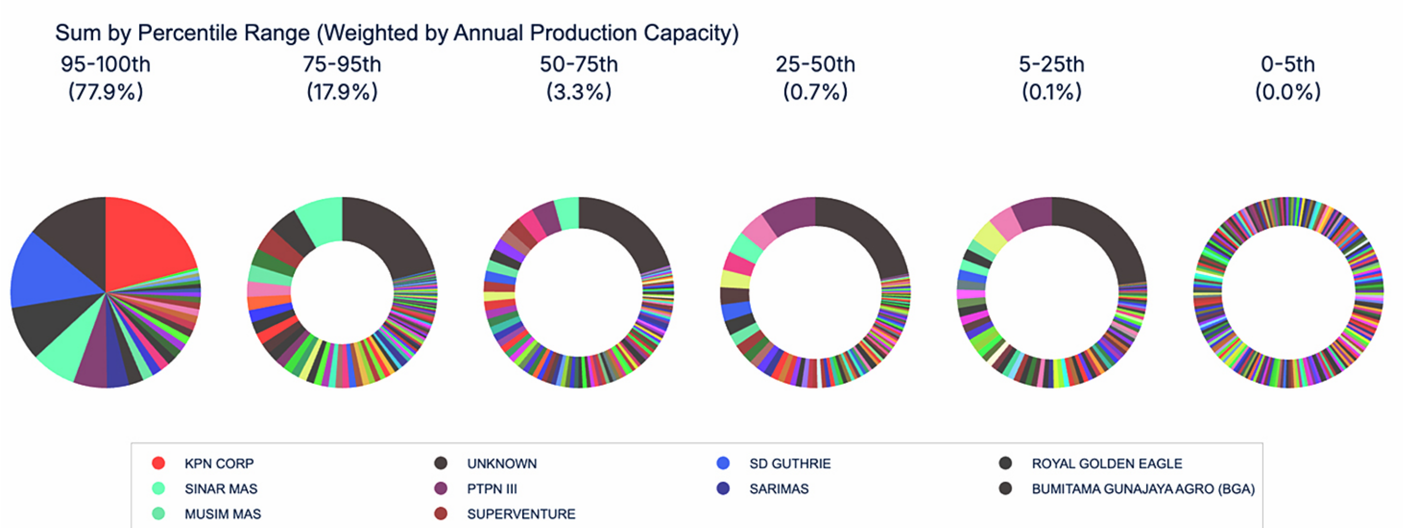

The order of the companies most exposed differs, but the top 5 remain the same. The reason for this change in order can be explained by the fact that some companies operate facilities for which deforestation over peat soil took place within their supply sheds. This has a disproportionately large associated Soil Organic Carbon (SOC) emissions loss, and could have pushed KPN corp to the top, with a 20.8% overall LUC emissions exposure (637,090 tCO2e/year).

4.4.2 Non-LUC Emissions

Non-LUC emissions are mostly driven by the amount of commodity growing area per supply shed. North Sumatra and Central Kalimantan both have clusters of relatively higher non-LUC emissions due to the higher density of commodity growing areas in those regions, but also due to the relatively larger supply shed size (see 4.5.1) for those facilities. This is mostly driven by the flat topography and good road infrastructure.

Figure 19: Per-supply shed LUC Emissions (tCO2e/year)

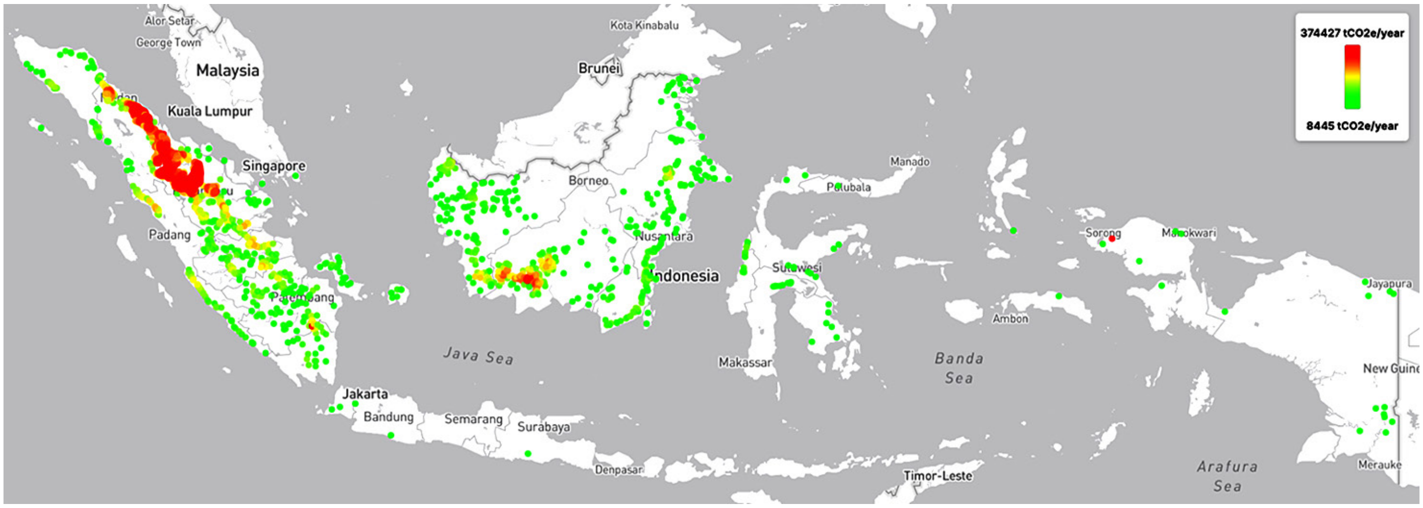

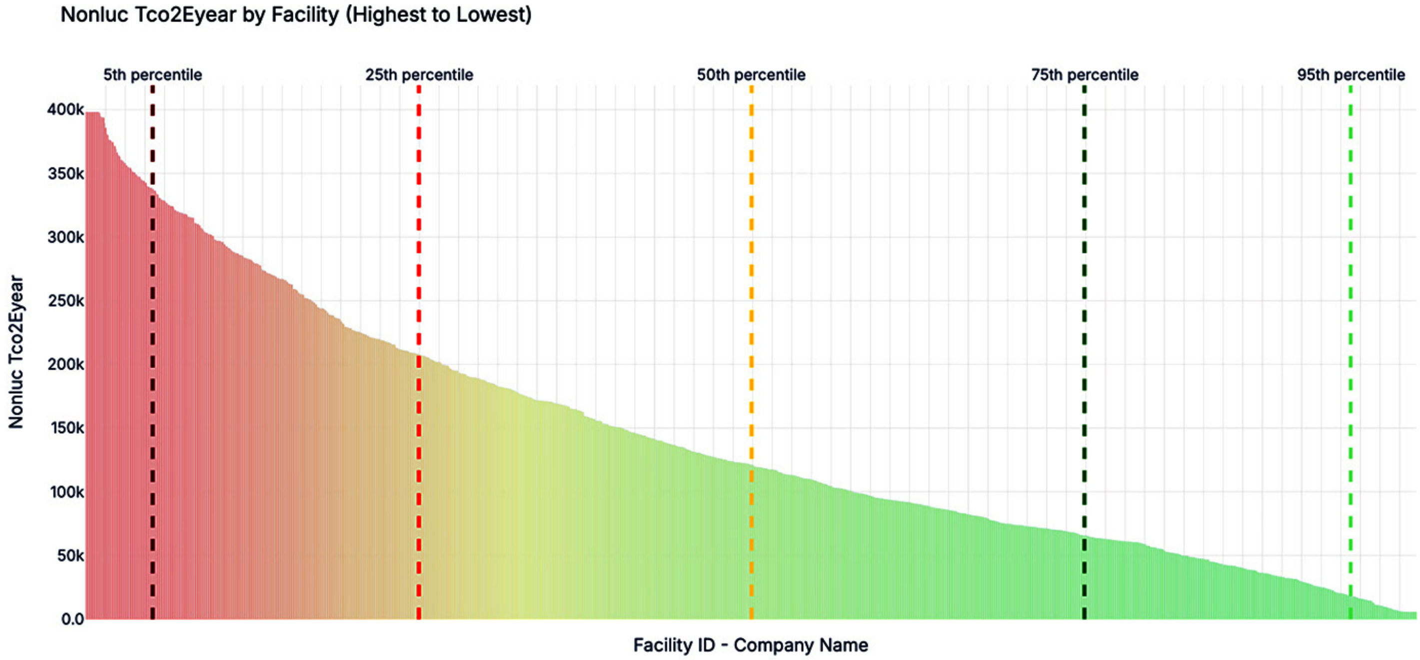

4.4.3 Biogenic Emissions

Biogenic emissions, which result from the clearing and growing of palm plantations, were intentionally reported separately, since they net out over the lifecycle of a palm plantation and are not incorporated in the overall environmental assessment of these facilities and their supply sheds.

Figure 20 shows that biogenic emissions may either be positive or negative. Positive emissions indicate that the supply shed has cleared more plantations than it has regrown in the monitoring period (2017-2024), whereas negative emissions mean more has regrown than cleared.

The biogenic emissions pattern indicates that plantation management is very sectorized, with a cluster in Northern Sumatra showing a high clearing rate of palm plantations having reached maturity, while other areas like northern Riau and Central Kalimantan have cleared very limited hectarage in the monitoring period (2017-2024).

Figure 20: Per-supply shed Biogenic Emissions by commodity growing area (tCO2e/year)

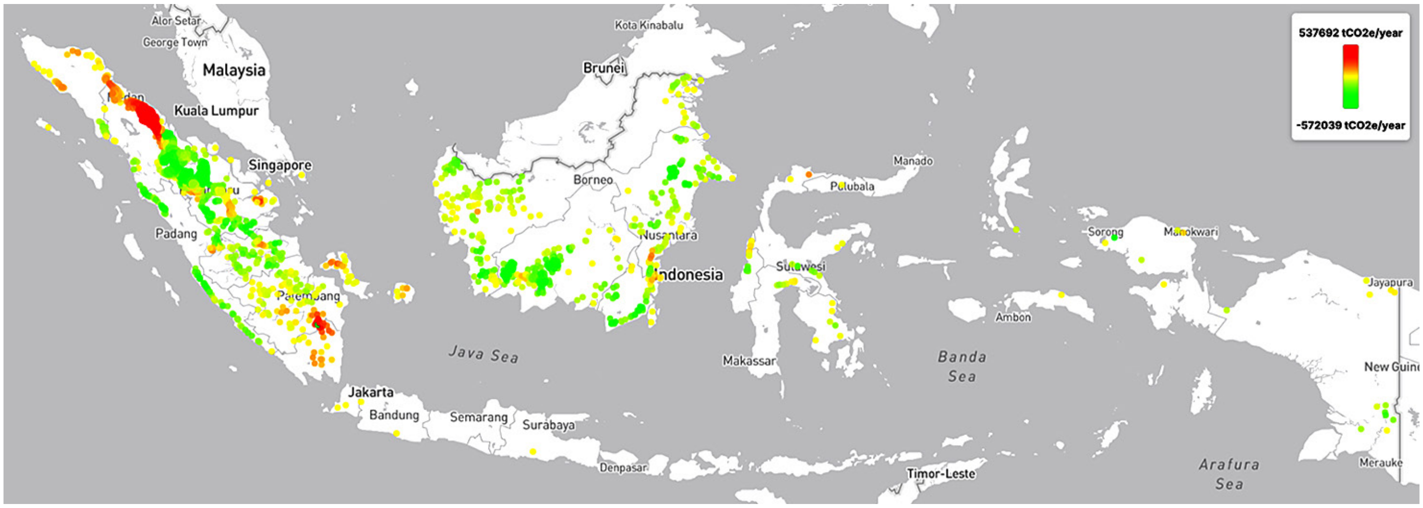

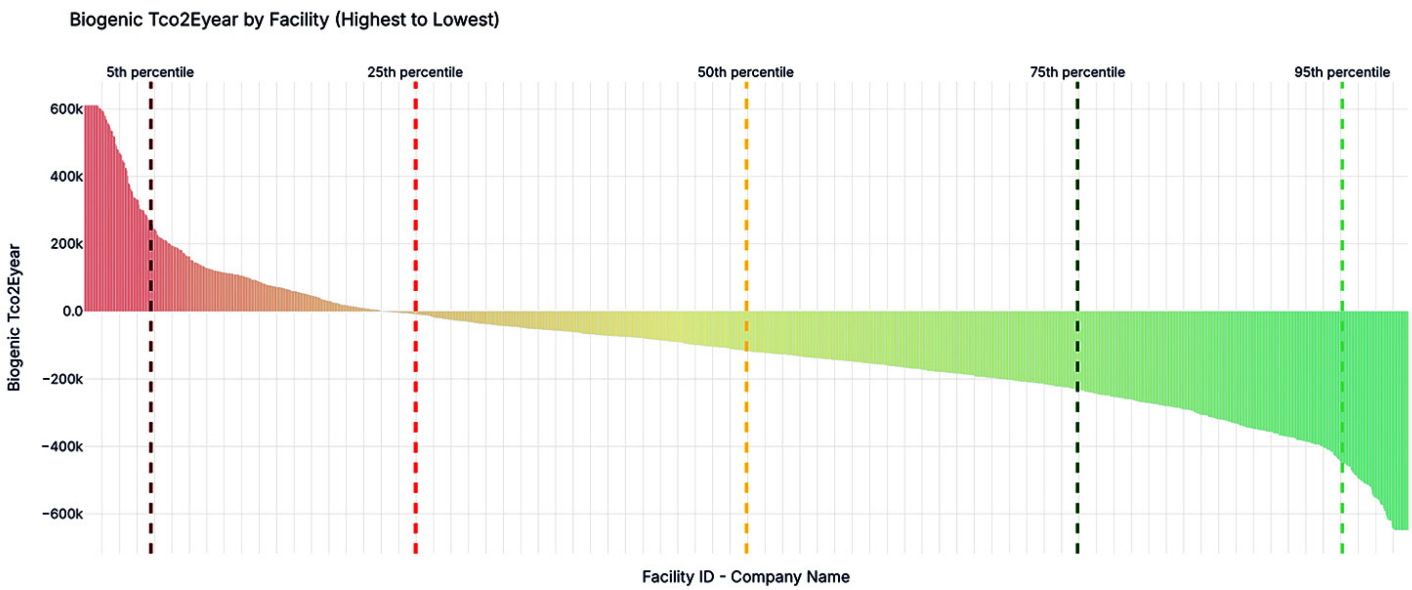

Figure 21: Relative contribution of Biogenic emissions per percentile range and per company. The size of the hole in the pie chart determines the relative importance of that pie chart for Biogenic emissions, with the first 2 pie charts representing "positive" emissions, and the last 4 pie charts representing "negative" emissions contributions.

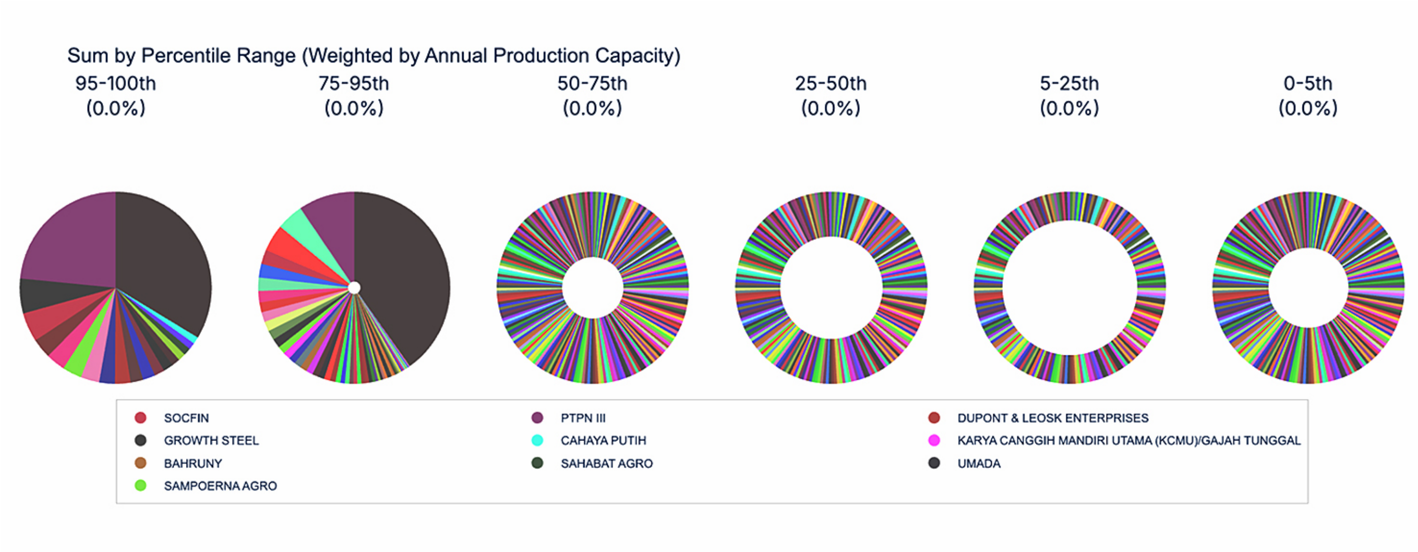

PTPN III has exposure to 23.5% of all biogenic emissions, thus indicating that they have been involved in high levels of plantation maintenance in the monitoring period (2017-2024).

4.5 Biodiversity

While canopy height heterogeneity is not a comprehensive metric of biodiversity, it does indicate the prevalence of monocultural practice (i.e. lower diversity score). Large parts of the dense estate areas of Sumatra and Kalimantan exhibit low diversity scores, while the smaller producing sectors like the south-western Sumatran coast and Sulawesi exhibit a higher diversity score.

Figure 22: Canopy Height Heterogeneity score ranging from 0 (red) to 1 (green)

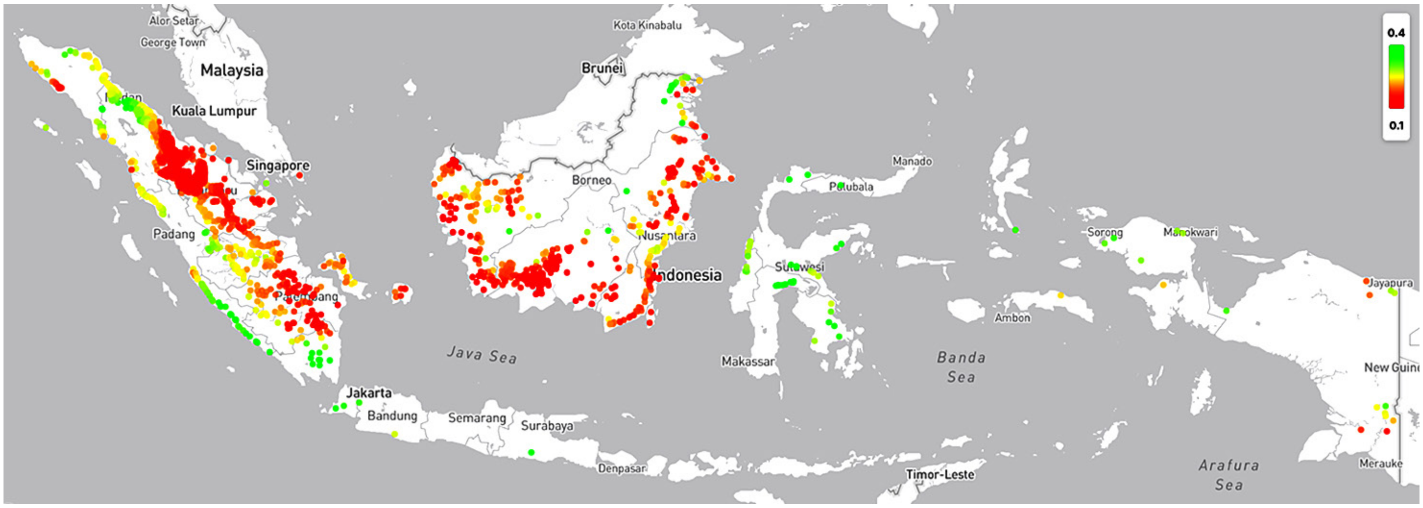

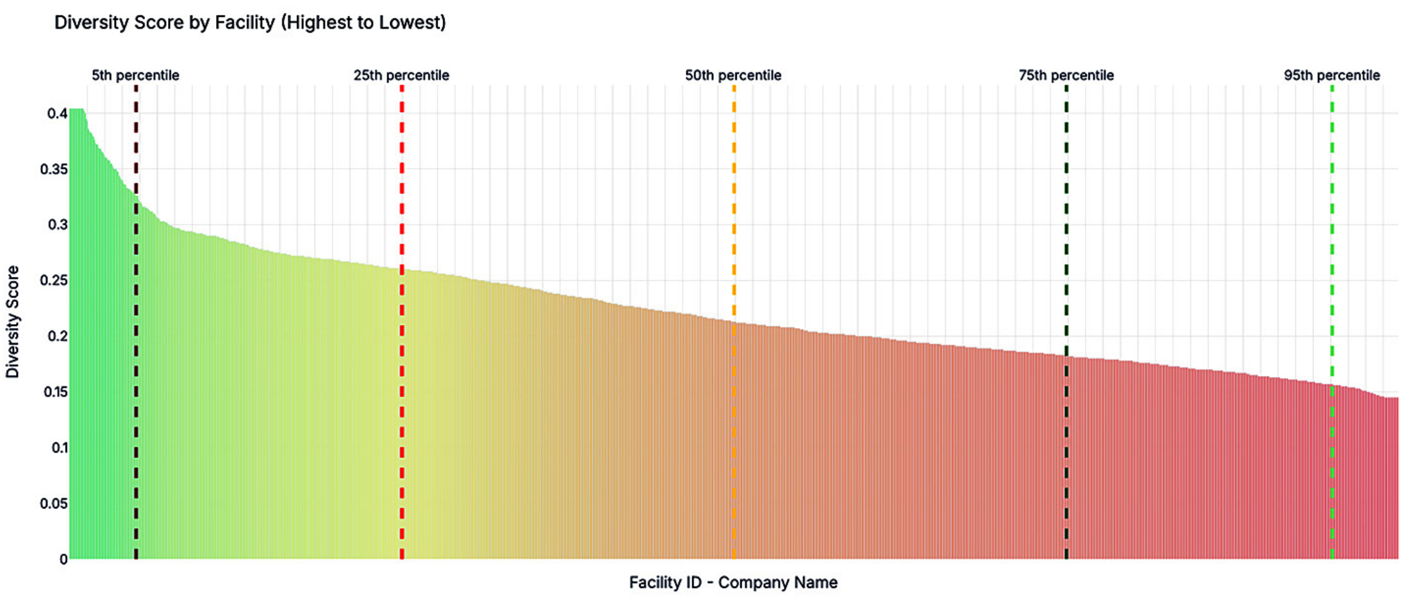

Figure 23: Comparison between one of the highest diversity score supply sheds (left: 0.39) in northern Sulawesi, and one with one of the lowest scores (right: 0.15) in Northern Riau

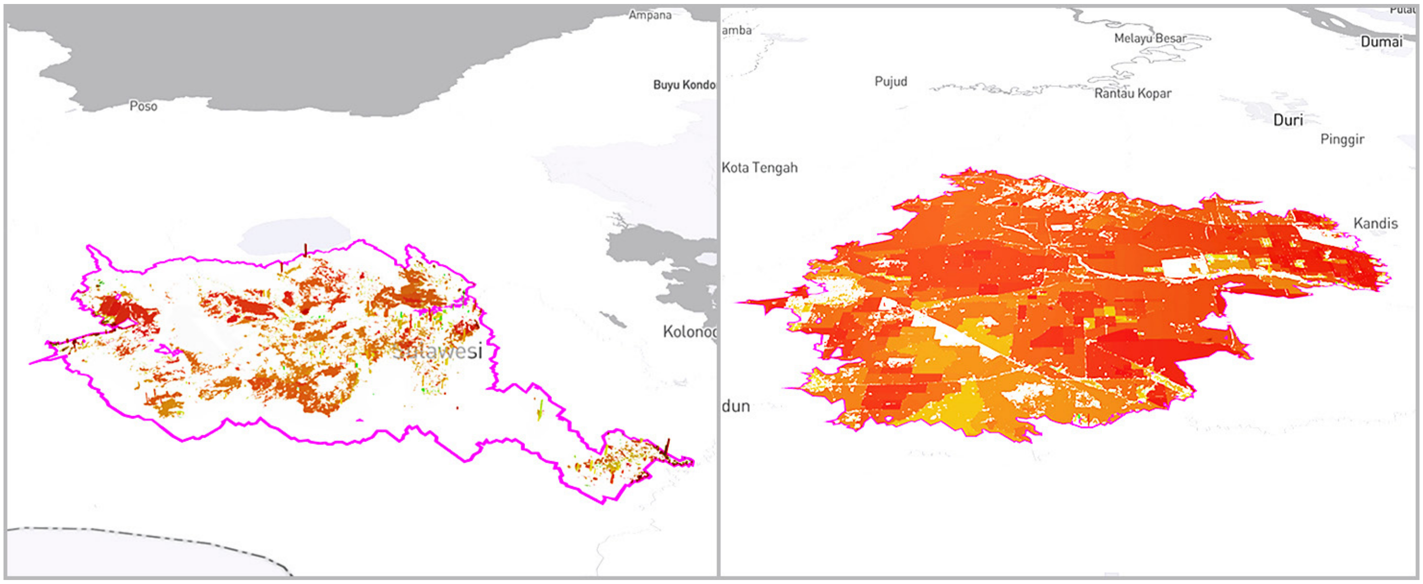

As illustrated in Figure 23, canopy height heterogeneity changes are mostly driven by the density and fragmentation of the plantations, with fewer smaller and sparser plantations leading to a higher diversity score, while the denser plantations lead to a lower diversity score.

4.6 Water Stress

Precipitation to potential evapotranspiration (PET) ratio seems to be low for the bigger part of the Eastern Sumatra and Kalimantan seaboard. This is unsurprising, as those areas are in the rain shadow of the Barisan mountains in Sumatra, and the Meratus mountains and the highlands of east Kalimantan for Kalimantan. In general, it is unsurprising that many areas exhibit a low precipitation-PET ratio, even in areas of abundant precipitation, due to the fact that supply sheds with high commodity growing area density typically experience higher air and soil temperature, which could exacerbate evapotranspiration, especially when comparing it to a natural forest ecosystem baseline.

Figure 24: Precipitation-PET Ratio ranging from 0 (red) to 3+ (green)

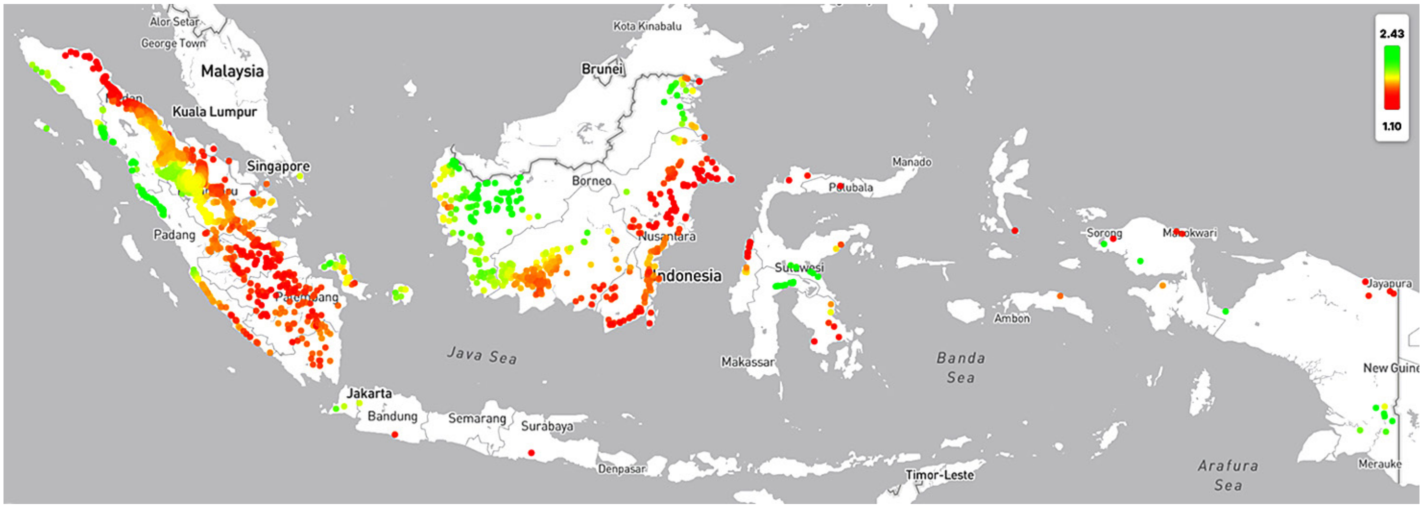

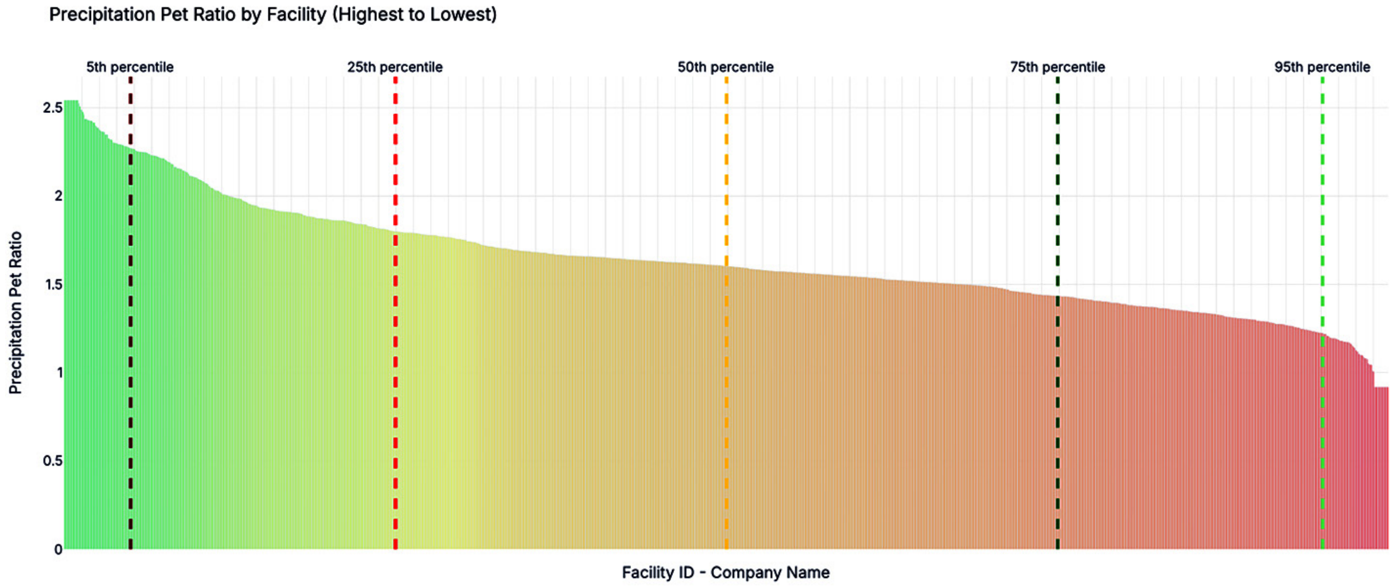

In terms of soil moisture percentile, the lowest values have identified the areas north-west of Medan (North Sumatra and Aceh's eastern seaboard), as well as areas in Eastern Kalimantan (West of Bontang). Two factors could explain the situation:

Dominant Soil Types with Lower Water-Holding Capacity: The predominant soil types in these upland and non-peatland areas are often Ultisols. These are ancient, highly weathered soils that, while suitable for oil palm with proper management, generally have a lower water-holding capacity compared to the vast peatlands (Histosols) found in other major palm oil regions like Riau and Central Kalimantan. Peat soils, when not excessively drained, act like a sponge, retaining significant amounts of moisture. Curiously, the areas identified to have a lower soil moisture percentile in Eastern Kalimantan and directly north-west of Medan are areas dominated by peat soils. This could indicate that the peat was drained in the course of the monitoring period (2017-2024), in which case the soil moisture percentile would be low when compared to the previous 5-year window.

Topography and Runoff: Many of the palm oil plantations in these specific areas are located on undulating or sloping land, especially in Aceh. This topography, especially on soils with moderate to low permeability, can lead to increased surface water runoff during rainfall events, reducing the amount of water that infiltrates and is stored in the soil profile.

Figure 25: Soil Moisture Percentile ranging from 0 (red) to 100 (green)

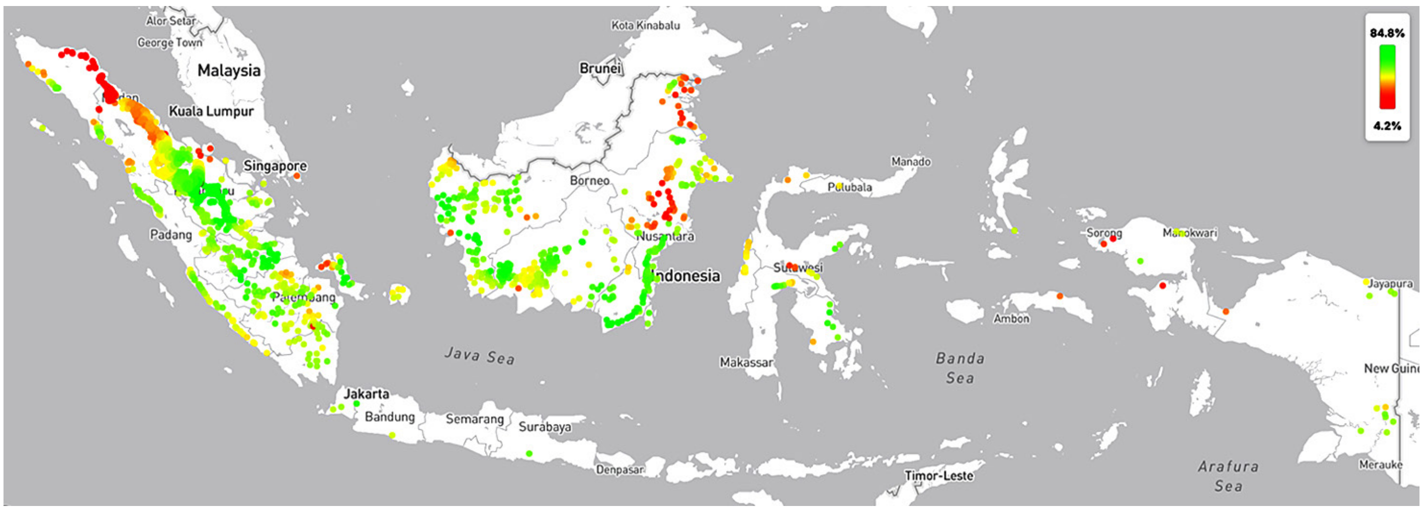

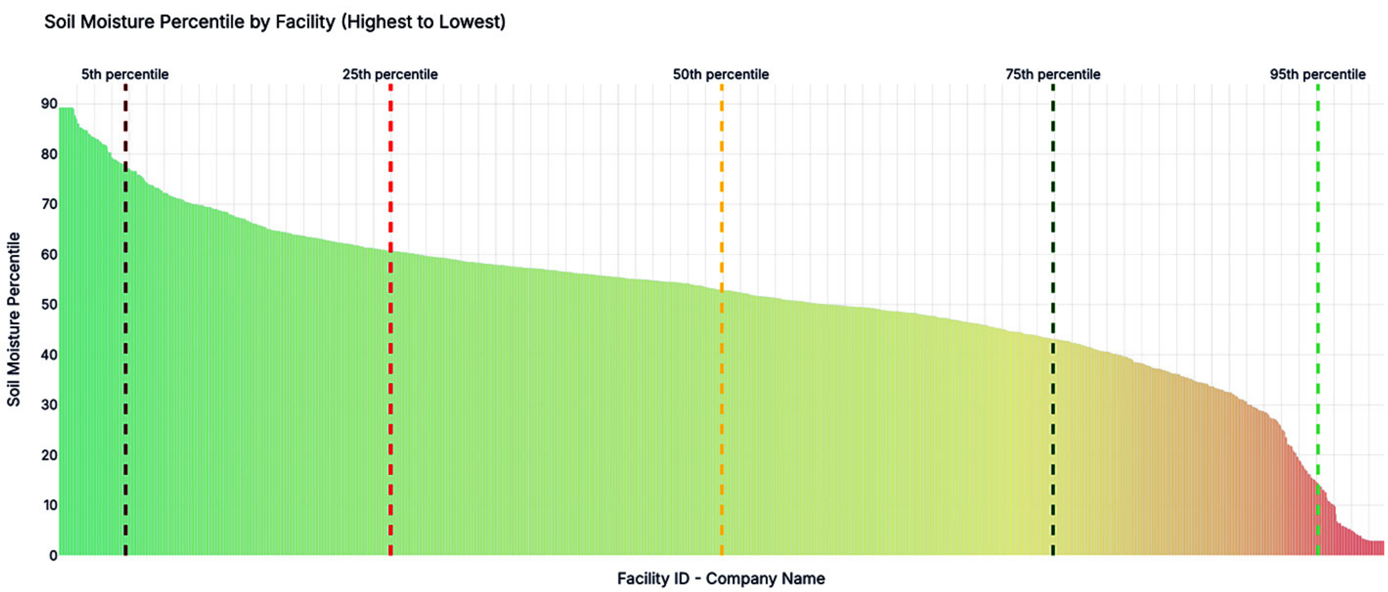

Actual evapotranspiration **anomaly** is a deviation from a 5-year preceding period average. Therefore, a negative evapotranspiration anomaly means that the amount of water evaporating from the land surface and transpiring from plants is **lower than usual** for that particular time and place. We see that big parts of Northern Riau, North Sumatra and West Kalimantan seem to have a positive actual evapotranspiration, which means that evapotranspiration is intensifying. This could be a synonym of both healthy ecosystems, or of commodity intensification. In the context of palm mill supply sheds, the latter scenario is more likely. However, even under the scenario of intensification, it still indicates proper ecosystem functioning, i.e. plantations are thriving, which could be thanks to irrigation or other measures to ensure plantation productivity. In the case of negative evapotranspiration anomalies, this could mean poor ecosystem functioning. This sub-indicator is harder to interpret than the others, but a negative anomaly is still considered to be worse than a positive anomaly. This sub-indicator only contributes to 20% of the final water stress weighted score for this reason.

Figure 26: Evapotranspiration anomaly ranging -1 (red) to 1 (green)

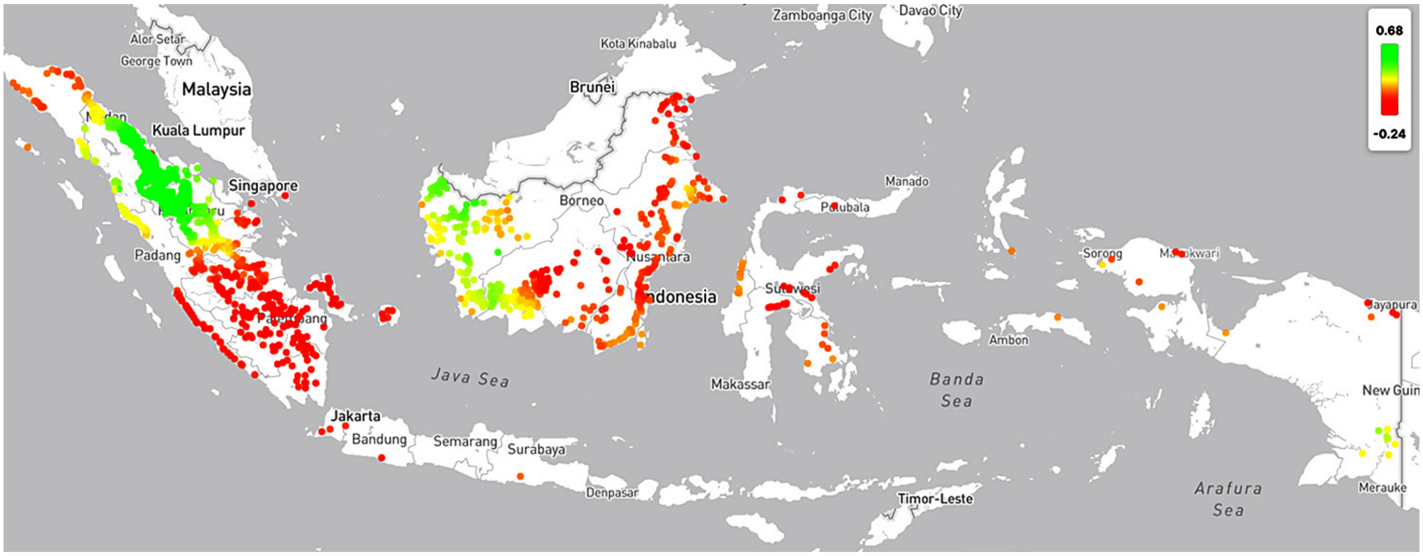

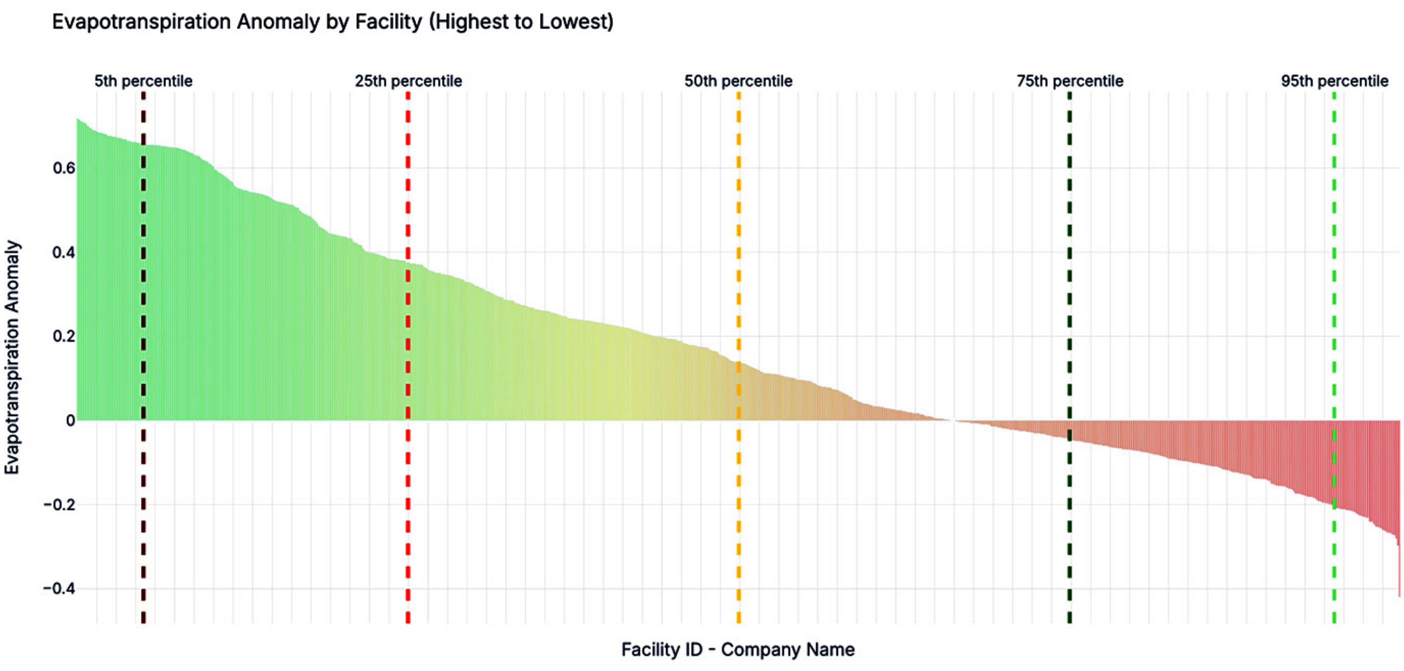

Based on the weighted water stress score, the areas coming out as most stressed seem to be the drained peat areas north-west of Medan and in Eastern Kalimantan (west of Bontang), and the areas on the steep slopes of the eastern Aceh seaboard. While this indicator disregards any hydrological flow in the supply sheds themselves and only provides a mostly meteorological perspective on the water stress problem, it has identified potentially problematic locations based on soil moisture deficit, which could be driven by the pattern of hydrological flow. Epoch expects to broaden its water stress assessment with supply shed-level hydrological flow modelling to provide a granular assessment of the impact of water at the plot level.

Figure 27: Weighted water stress score ranging from 0 (red) to 1 (green)

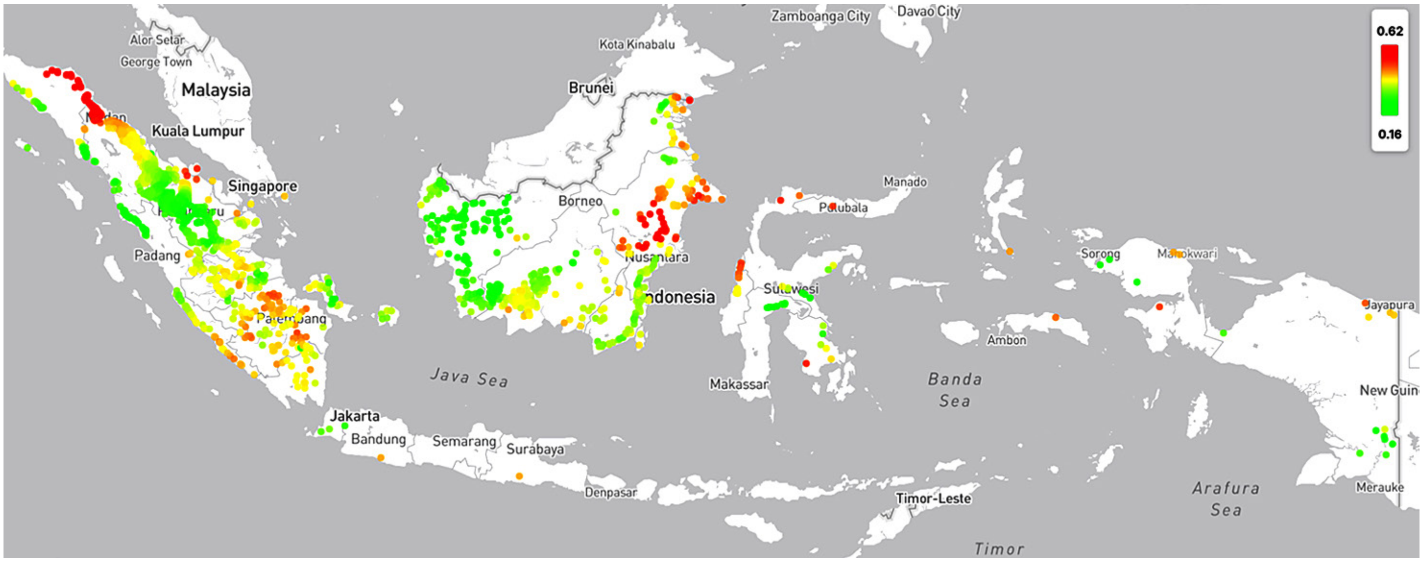

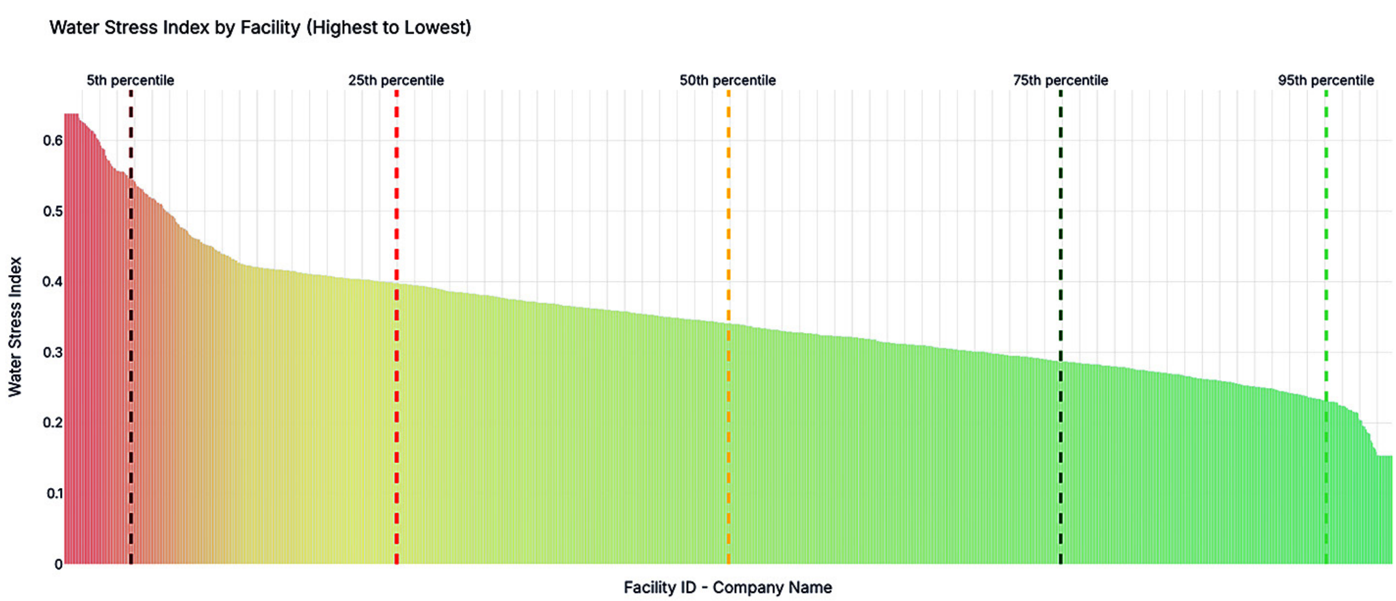

Figure 28 and Figure 29 respectively show the landscapes where water stress was identified, the former due to plantation on steep slopes that limit the ability for soil moisture to be retained, and the latter peat drainage to expand oil palm estates in Eastern Kalimantan which creates a sudden drop in actual evapotranspiration due to lower levels of soil moisture and reduction in above-ground biomass.

Figure 28: Small- and medium-scale palm plantations on steep slopes in the south-east of the Bireuen Regency, Aceh

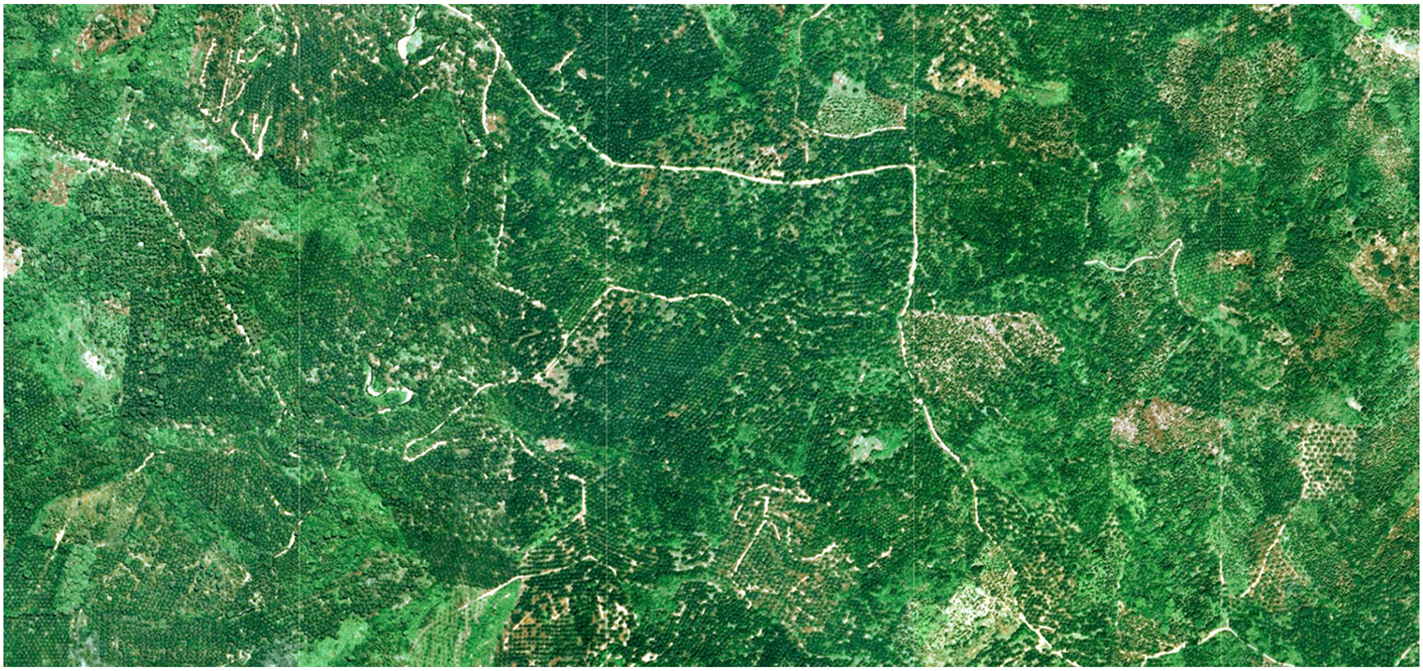

Figure 29: Small- and medium-scale palm plantations on steep slopes in the south-east of the Bireuen Regency, Aceh

4.7 Overall Risk Score

If we attempt to weigh all inputs above into a single risk metric using percentile ranking (i.e. this purely provides a relative risk between facilities, and not an absolute risk score), the facilities most at risk expectedly appear around the intensified estate areas of Central Kalimantan and the northern Riau and North Sumatra. Those areas are the most intensified palm production areas in Indonesia, with highest rates of estates and lowest rates of smallholders, and exhibit the largest contiguous areas of oil palm production across all palm growing regions in the country.

Figure 30: Overall risk score ranging from 0 (red) to 1 (green)

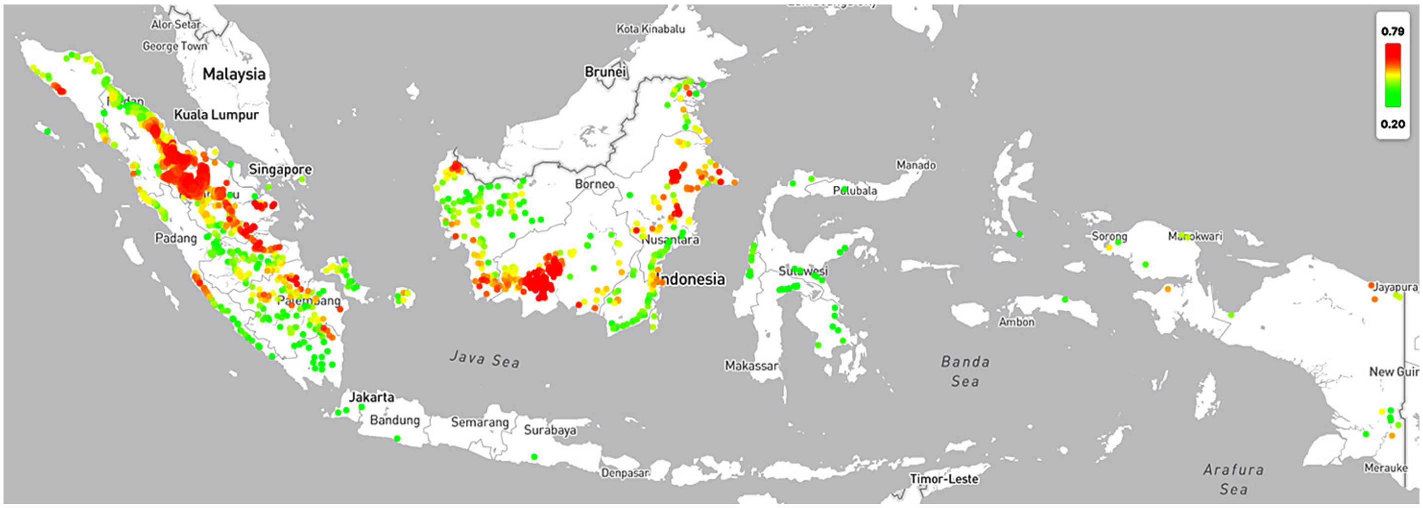

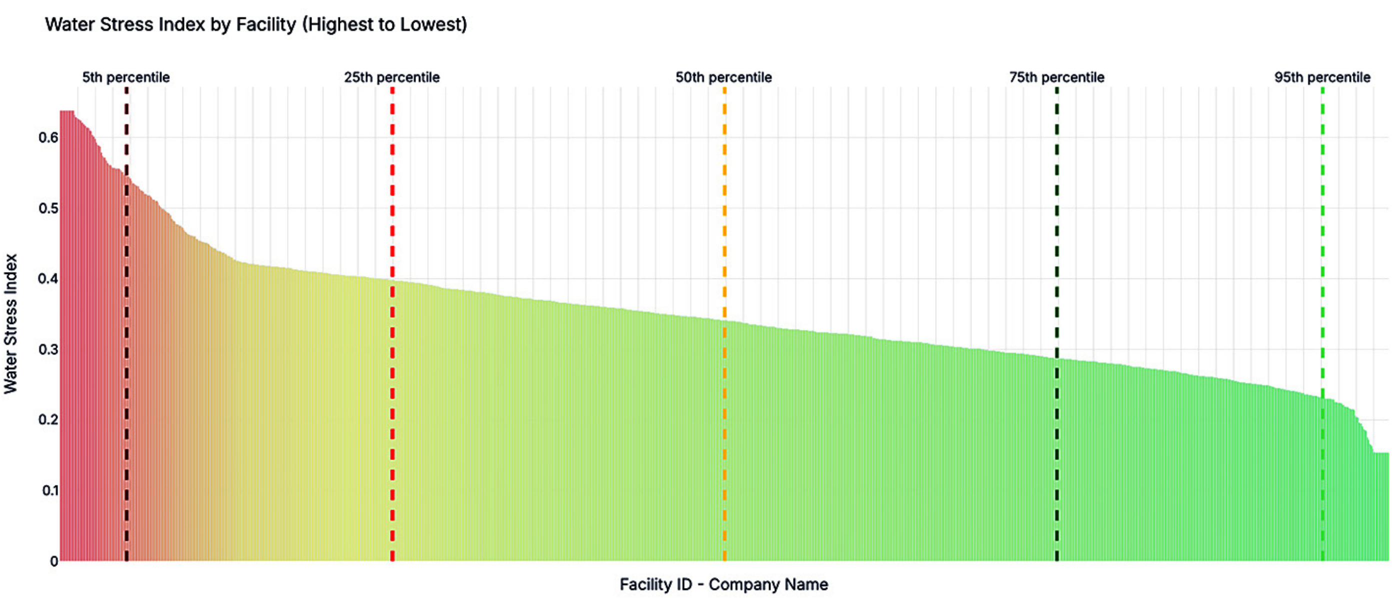

4.8 Relative Contribution for the “Bottom Performers”

For each environmental indicator, we look at the relative contribution of the 0-5th and 5-25th percentile bins of producers and report the corresponding production capacity affected by the corresponding percentile bin.

Table 9: Relative contribution of the 0-5th percentile for the 4 main metrics produced. The last column represents the percentage of the entire palm production capacity for Indonesia.

Metric	0-5th Percentile Absolute Contribution	0-5th Percentile Relative Contribution	0-5th Percentile Absolute Annual Production Capacity	0-5th Percentile Relative Annual Production Capacity
LUC Emissions	3,064,899 tCO2e/year	77.8%	19,200,000 ton FFB	5.2%
Deforestation	140,685 ha	73.2%	20,310,000 ton FFB	5.5%
Water Stress	0.61	27.8%	12,318,000 ton FFB	3.5%
Biodiversity	0.01	10.5%	20,640,000 ton FFB	5.9%

Table 10: Relative contribution of the 5-25th percentile for the 4 main metrics produced, and their corresponding production capacity. The last column represents the percentage of the entire palm production capacity for Indonesia.

Metric	5-25th Percentile Absolute Contribution	5-25th Percentile Relative Contribution	5-25th Percentile Absolute Annual Production Capacity	5-25th Percentile Relative Annual Production Capacity
LUC Emissions	712,002 tCO2e/year	77.8%	75,678,000 ton FFB	20.6%
Deforestation	41,815 ha	21.8%	74,798,000 ton FFB	20.3%
Water Stress	0.44	20%	62,604,000 ton FFB	17.9%
Biodiversity	0.02	12.1%	75,750,000 ton FFB	21.7%

The affected percentage of production capacity seems to lie around 5% for the 0-5th percentile and 20% for the 5-25th percentile respectively, which is proportional to the relative production volume. Only water stress represents a smaller % of overall production capacity.

5 Discussion

This pilot was designed to answer three core questions for the *Codex Planetarius* initiative: how to identify the producers responsible for a disproportionate share of environmental harm, how to translate farm-level impacts to their first point of aggregation, and how to establish a scalable methodology for doing so. The results from this Indonesian palm oil analysis provide a clear and affirmative answer to each. The discussion below explores the key trends revealed by this methodology, evaluates the validity of its core assumptions, and outlines a clear path for improvement and expansion.

5.1 Trends in the Data

The supply shed analysis reveals a stark concentration of environmental risk within the Indonesian palm oil sector, challenging the notion of a uniformly distributed problem and highlighting the critical importance of targeted interventions. The findings demonstrate that the most severe impacts, namely deforestation and associated LUC emissions, are not widespread but are overwhelmingly linked to a small fraction of the supply chain; the bottom 5% of facilities are associated with over 73% of all post-2020 deforestation detected in this study. This disproportionate impact is further concentrated among a handful of corporate actors operating primarily in southern Riau and Jambi. Moreover, the overall risk score indicates that the most intensified, estate-dominated supply sheds in Central Kalimantan and northern Sumatra, while not necessarily driving recent deforestation, exhibit the lowest biodiversity scores and significant water stress. This suggests a critical trade-off where highly productive monocultures become hotspots for other environmental vulnerabilities, reinforcing the value of this supply-shed-level approach to pinpoint not just where problems exist, but precisely which actors and landscapes require the most urgent engagement to meet sustainability goals.

These findings validate the central premise of *Codex Planetarius*: that a targeted approach, focusing on the worst performers, is not only possible but is the most efficient path toward systemic change.

5.1.1 Supply Shed Approach – A Unique Data Aggregation Approach

Supply Shed delineation allows adequate attribution of environmental metrics to the associated facility, in a way that is much more representative than a jurisdictional assessment would. Jurisdictional assessment may support governments in reporting emissions or making evidence-based policy, but for as long the data is not facility/supplier specific, the use of this type of information will remain limited. Producing this information at supply shed level with company ownership associated allows deeper use cases like financing (multilateral as well as private), insuring, and targeted investments for landscape interventions. If the environmental credentials cannot be tied to a specific value chain, it is virtually impossible to engage with the right actors and the right sourcing landscapes, and the problem remains a tragedy of the commons. The difficulty of this approach is to enrich the facilities and supply sheds with company ownership information, which allows tying back the plot and supply-shed level information produced back up the value chain. The Trase approach to gather this information should be studied, and see if the playbook can be reproduced systematically for every new commodity and country for which this supply shed methodology is applied. Guye, 2024 and Mempel et al., 2025 have carried out similar foundational field work to identify this information. With 3 experiments carried out for 3 different commodities and 3 different countries, gathering lessons learnt from the logistical and budgetary challenges of doing this data collection and validation at scale should be further elaborated.

5.1.2 Deforestation and LUC Emissions – Geographically Concentrated

When looking at the deforestation and LUC emissions results of the supply shed analysis, the exposure is extremely skewed towards few actors. The major pitfall of this analysis is that only a 2017-2024 monitoring period is taken, and it is known that in palm oil most of the land conversion took place more than a decade ago for the majority of the established estates. From an emissions accounting perspective, reporting the LUC emissions figures in this report as part of a SBTi report would not be aligned with the GHGP LSRG guidance. However, we are mainly interested in forward-looking use cases here, namely understanding the patterns of land use within the supply sheds in the recent past and supporting legislations and investments to mitigate the natural land conversion challenge.

5.1.3 Estate-Dominated Supply Sheds Most Exposed to Risk

When looking at the overall results of the supply shed analysis, it is clear that supply sheds with a high proportion of estate areas tend to exhibit lower environmental scores. While the direct causal link between established, older estates and deforestation or land-use change (LUC) emissions may be less pronounced, these extensive plantations nonetheless exert considerable pressure on local ecosystems. Specifically, they demonstrably impact water availability through increased demand and altered hydrological cycles, and contribute to significant losses in biodiversity due to habitat fragmentation and monoculture dominance. Additionally, on sloped terrain, the lack of ground cover and intensive cultivation practices can lead to significant soil erosion, diminishing soil fertility and contributing to sedimentation in waterways. Furthermore, these large-scale operations exacerbate environmental risk factors, notably increasing the likelihood of pest and disease outbreaks, and elevating fire hazards, often linked to the management practices within these estates and their proximity to vulnerable ecosystems.

5.2 Validity of Assumptions and Potential Improvements

5.2.1 Completeness and Accuracy of Trase Palm Oil Facilities

To our knowledge, the facilities dataset used as part of this study is the most comprehensive that exists for the palm oil industry in Indonesia. The screening of the facility locations reveal high data integrity from a precision point of view. There is still a chance that some newer facilities (from 2024 onwards) may not have been identified in the Trase dataset. Routine identification of new palm oil facilities is entirely considerable in the age of geospatial embeddings availability. The potential to reduce effort and cost to identify and enrich a palm oil facilities database using embeddings similarity search is enormous, and Epoch is currently working on models to do this, as Figure 31 illustrates.

Even though some facilities may be missing from the Trase dataset, it does seem like the dataset is both highly accurate in terms of coordinate location, and near-complete in terms of coverage as of 2024. It is therefore an excellent dataset to use to power the analysis presented in this report.

Figure 31: Example of a facility (1,7177° N, 100,217315° E) missing from the Trase Dataset, generated from a similarity search in the AlphaEarth embeddings dataset for 2024 from Google. The Google embeddings incorporate a spatial attention mechanism in a 3x3 convolution (Brown et al., 2025), which reinforces the detection of certain features in a given spatial context (i.e., industrial facility + ponds + palm plantation), enabling the retrieval of such facilities reliably (and why some of the surrounding plantation has been included in the detection).

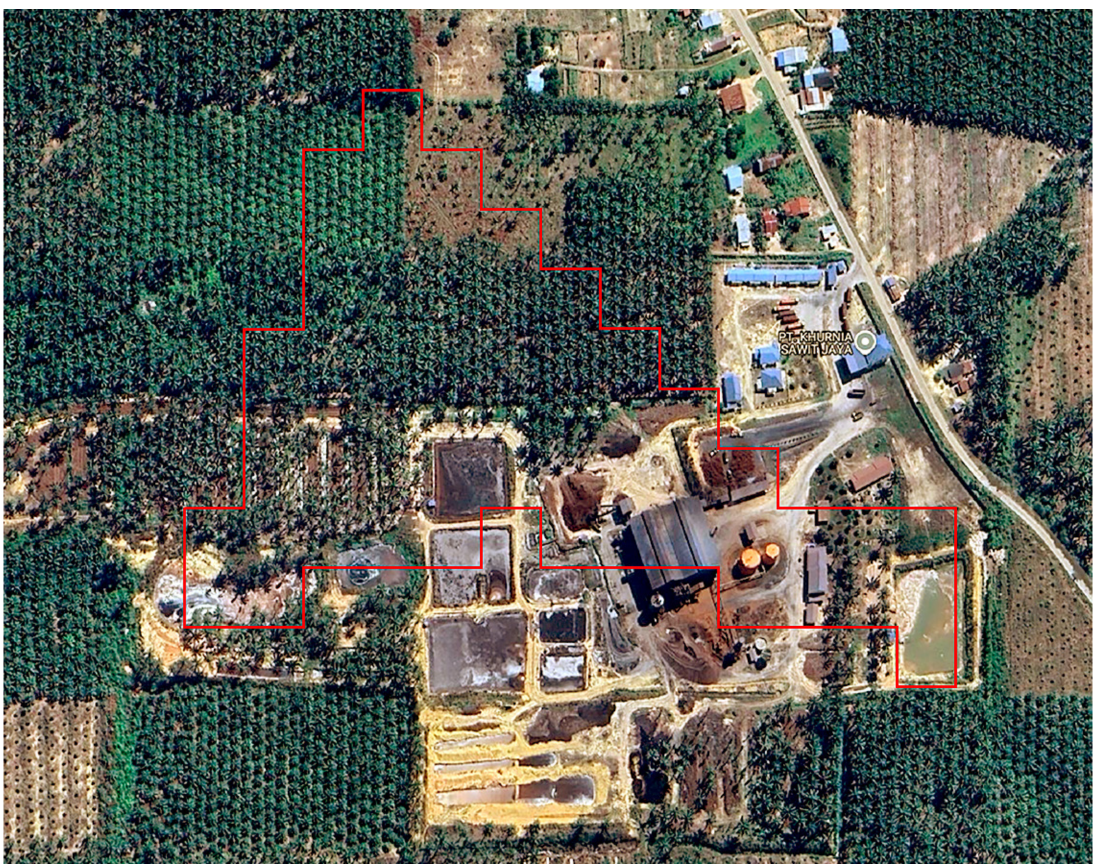

5.2.2 Representativeness of the Generated Supply Sheds

The major pitfall of this approach is the difficulty to validate the representativeness of the derived supply sheds, for two main reasons:

- **The supply shed may under-capture the entirety of the potential palm sourcing area.** Some areas may remain out of the reach of the derived supply shed, even though it should belong to it. A future improvement could be to ensure that any palm growing area is assigned to a facility (the closest one by road distance), even if it lies outside of any supply shed boundary. This would ensure that no palm growing area is unassigned.

Figure 32: A supply shed in Maluku falling short of capturing additional palm plantations, which necessarily belong to the catchment of the facility, since it is the only one on the Island

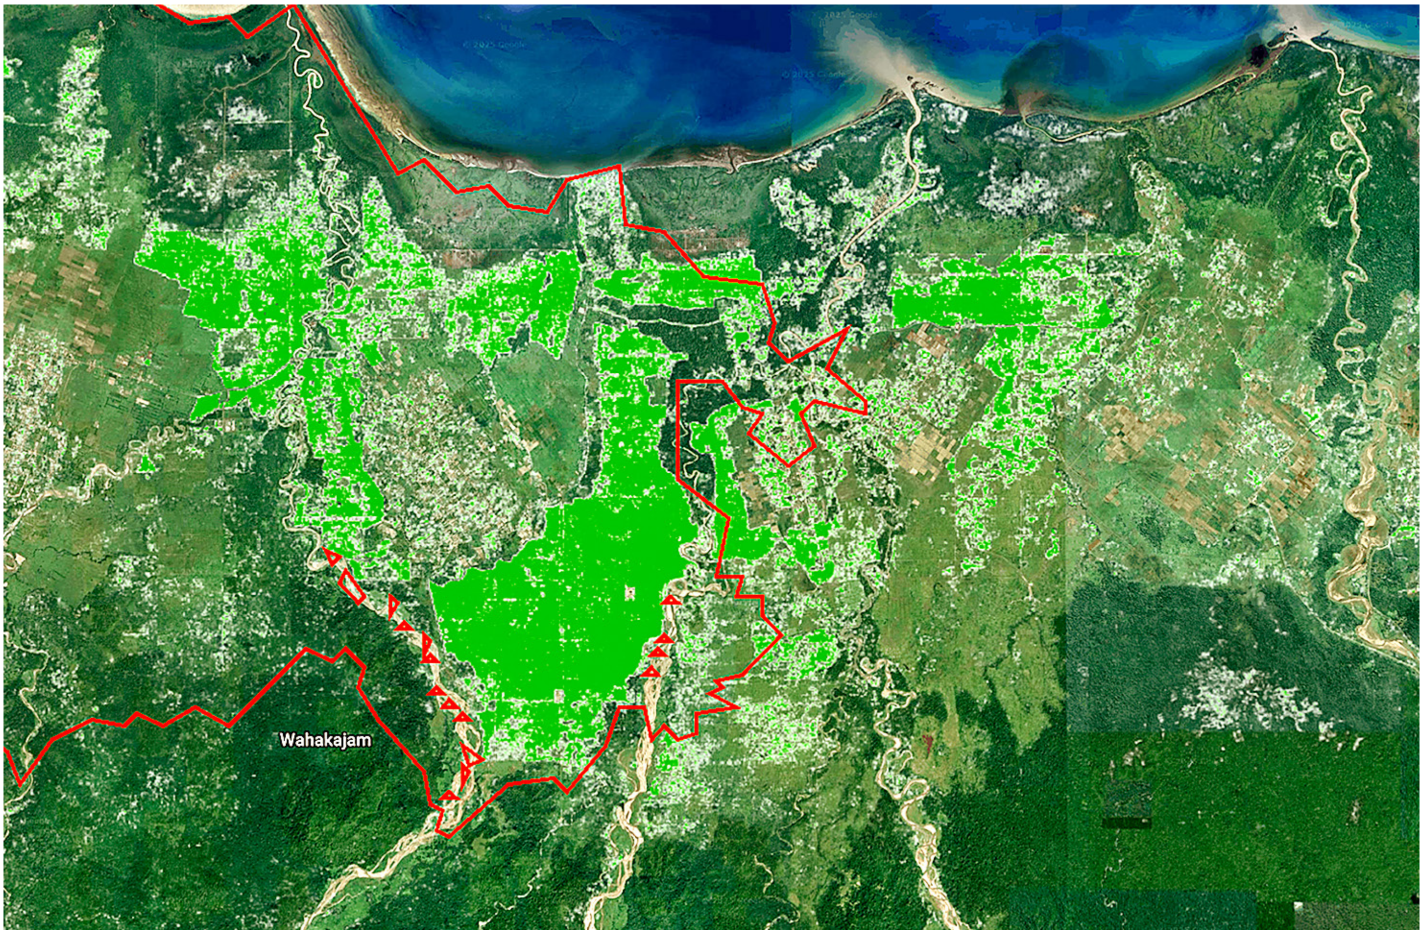

- **The supply shed may over-capture the potential palm sourcing area.** Due to the probabilistic nature of the approach, a supply shed should be conservative (i.e. intentionally larger) in order to avoid underestimating the environmental exposure of a given facility. However, in an area of high facilities density, supply sheds will inevitably overlap, making the probabilistic assignment of overlap areas difficult. Although not implemented this time around, this would require an elaborate allocation approach to intelligently allocate each palm growing plot (or pixel) to an appropriate supply shed based on the inter-supply shed least cost path comparison. The naive way would be to assign the palm plot/pixel fully to the supply shed for which the least cost path calculated was the shortest. The more elaborate way would be to assign each pixel/plot in all the supply sheds with a probability of belonging to any given supply shed. The latter option would require more elaborate computations but would probably be more representative of the fruit bunch off-taking reality.

Figure 33: An example of supply sheds overlapping in Central Sulawesi. Considering the level of overlap in some of their cases, a probabilistic allocation based on least cost path would be most appropriate, so that the environmental metrics produced for the respective supply sheds can be weighted using this probability of belonging to either of the supply sheds.

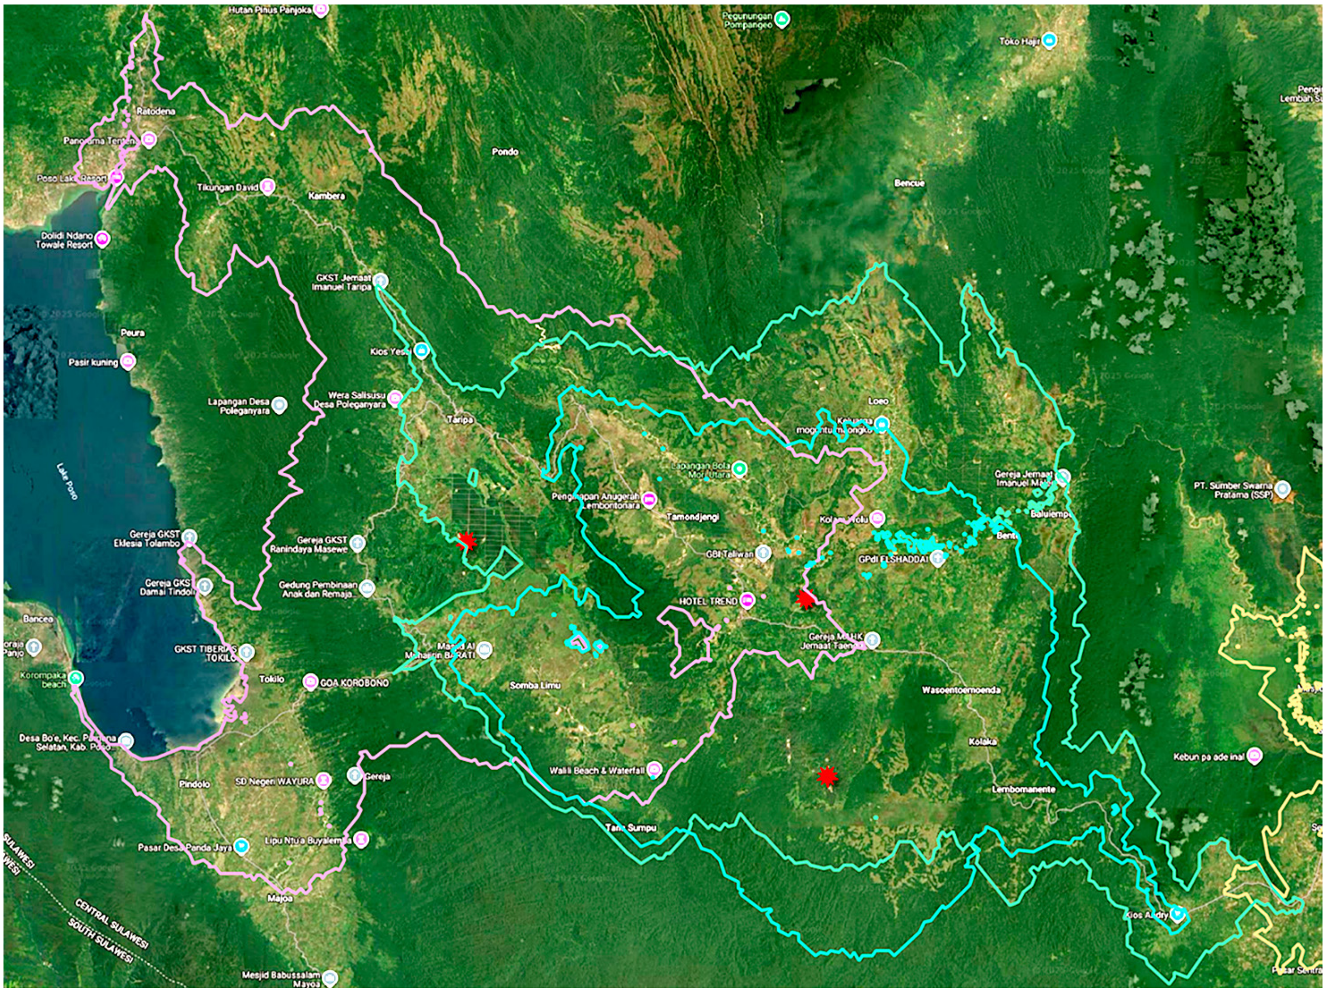

Moreover, the production capacity metadata field in the Trase facility dataset was not leveraged in the supply shed calculation consciously, assumption being that a smaller facility can still technically source produce from equally far from a larger facility, because the freshness requirement does not change. However, what could have been done is to use the production capacity as a weight for the aggregated environmental metrics, so as to reflect the true potential impact of the facility within its supply shed. An appropriate way to do this will be considered in a new iteration.

The concession areas used to refine the supply shed boundaries from Global Forest Watch date back to 2016, and having access to a more up-to-date pan-Indonesian dataset would greatly improve the supply shed outlines.

Finally, adopting a machine learning-based approach to classifying smallholder vs estate extent similar to how the data produced in Descals et al., 2024 would greatly improve the delineation of the respective production areas, especially if using spatial embeddings as inputs to the process.

5.2.3 LUC Emissions – Model Over Long or Short Time Scales?

Industries like Oil Palm that have a uniquely large interface with nature due to the nature of their sheer size and the fact that they grow in some of the most biodiverse and ecologically rich ecosystems on Earth. From that perspective, agricultural practices that convert natural land, especially peatland forests, even if they took place a long time ago, are important for an emissions accounting and environmental footprint perspective. In that regard, the fact that this study looked at the 2017-2024 epoch without considering prior information is reductory and does disjustice to having an adequate overview of the impact of LUC emissions for the different facilities and the companies that own and operate them.

SOC emissions factors for mineral and peat soils were still applied to all areas with a plantation age < 20 years, since peat SOC emissions are disproportionately large, even if they took place a long time ago. Even though fairly accurate delineation of peat and mineral soils are available (PEATGRIDS, Widyastuti, 2024) to allow performing spatial allocation of the emissions factors, using emission factors remains a gross simplification. Most SOC and/or peat carbon stock data products are not change products and only achieve carbon stock estimates for a given snapshot in time, making it difficult to model SOC change adequately. A first approach could be to use mapped peat depth to determine the gravity of peatland forest conversion, rather than rely on a uniform emission factor. A more elaborate approach could be to the SoilDB dataset from Hengl et al., 2025, which maps temporal soil properties temporally for the period 2000-2022. It is unlikely the data models peat draining events and other abrupt land conversions accurately, but it may still prove to be more accurate than the current emission factor-based approach.

Moreover, attempting to harmonize modelling done with specific input types over a longer period of 20+ years, and modelling done over a more recent time period such as for this study, is non-trivial. Different input resolutions, modelling assumptions and biases would make the alignment of the results difficult, especially since the bulk of the emissions modelled depend on temporally-consistent biomass (and potentially soil organic carbon) stock retrieval information.

In general, *Codex Planetarius* seeks to move the needle on policy-making and financing going to the bottom 5-10% performers, making the aggregation and retrieval of legacy information less relevant to meet its objectives. That said, painting the picture which aligns with greenhouse gas accounting best practices (SBTi/GHGP LSRG) is a definite improvement that could be made to this methodology.

5.2.4 Non-LUC Emissions

Emissions around farmgate practices and processing are well documented for the palm oil value chain, as well as its variability based on peer-reviewed studies, but such data is not widely available across all palm producers, and will probably not be any time soon due to the cost and coordination required to collect this data in a comprehensive manner.

The main farmgate emissions sources that are not tied to land use change is fertilizer application (mainly the origin of the product and its quantity applied), which can truly only be known if the supplier is willing to share this information. This may happen on a case-by-case basis if a company is willing to disclose this information for formal GHG reporting purposes, but will not synchronously happen across all 1400+ facilities.

On the other hand, for emissions at processing level, a few features could be identified that correlate with facility emissions:

- **Palm Oil Mill Effluent (POME):** Wastewater effluents in palm are rich in organic matter. When stored in open ponds or lagoons, it undergoes anaerobic decomposition, releasing large quantities of methane (CH_4) and carbon dioxide (CO_2). The volume and characteristics of POME are directly related to the amount of FFB processed, which could help determine whether the facility is a high-, mid- or low-throughput installation. The exact cause-to-effect of number of ponds and their size cannot be determined due to varying factors like pond depth which cannot be readily retrieved however.
- **Methane capture systems (biogas domes) and associated storage tank capacity:** The presence of this additional feature can be a critical determinant of the POME emission factor. For a mill with open POME ponds, the emission factor is very **high**. A large amount of methane is emitted for every tonne of palm oil fruit processed. For a mill with a biogas dome, the emission factor is drastically **lower**, potentially near zero for POME. The primary greenhouse gas is captured and either converted to a less potent gas or used for energy.

The number and size of POME ponds could be automatically retrieved to infer the appropriate emissions factor to apply.

Alternatively, direct methane estimations of palm point sources could be considered by using satellite missions dedicated to this purpose (e.g. MethaneSAT) by combining methane concentration data with wind speed and direction, it is possible to estimate the rate of methane release from a specific source.

Table 11: Identifying facilities which have methane capture systems from those that do not has a significant impact on the emissions factor of the facility

Treatment Method	Emission Scenario	Emission Factor (kgCH4 / ton FFB)	Key Determinant	Peer-Reviewed Source Indication
Open Anaerobic Ponds or direct discharge to river	High Emissions	~6.54 kg CH4	No methane capture is in place, allowing direct release of biogas to the atmosphere.	Schuchardt et al. (2008)
Methane Capture System	Low Emissions	~1.31 kg CH4	A biogas dome or cover captures the methane. The emission factor reflects an ~80% capture efficiency.	Judijanto et al., (2025)

POME-related emissions is by far the largest processing emission source, representing typically more than 65% of all processing emissions, followed by electricity consumption at 25% and transport (<10%). For the latter two, connection to the electricity grid of a facility could be inferred by leveraging Overture Maps infrastructure data (power lines and substation locations), to determine whether there is even a chance of a facility being connected to the grid. For facilities in the vicinity of the physical grid infrastructure, VIIRS night lights could be analysed to check the brightness of the facilities across the year. Bright spots above a certain threshold would likely be connected to the grid. Finally, transport emissions could easily be inferred from the road distance from the farms to the facility, and from the facility to the closest refinery or port of export.

5.2.5 Biodiversity Modelling - Landscape Connectivity Missing from the Picture

Canopy height heterogeneity, a simple metric, indicates structural diversity within a plantation, moving away from monoculture risks. High heterogeneity suggests uneven-aged stands, shade trees, or agroforestry, indicating on-farm biodiversity. However, this metric's limitation is scale; it shows "what" but not "where" in the ecological network, failing to reveal proximity to crucial forest remnants or riparian corridors essential for resilience. It measures farm texture but ignores surrounding life-support systems. Therefore, while canopy height heterogeneity signals a move from monoculture at the plot level, it requires landscape-level analysis of habitat fragmentation and connectivity for true operational resilience.

Ecological science increasingly challenges the wisdom of vast, homogenous monocultures for resilience. Integrating natural buffers and fragmentation within a palm oil supply shed is a strategic imperative. Landscapes with natural habitats like forest patches, riparian corridors, and hedgerows mitigate pest and disease outbreaks, safeguarding the supply chain.

This resilience stems from ecosystem services provided by natural buffers. Studies show landscape complexity enhances pollinator and pest predator populations. In palm oil, fragmented habitats serve as refuges for natural enemies of pests (e.g., parasitic wasps), reducing reliance on chemical pesticides. This "biological control" diminishes in large monocultures, leading to pest explosions.

Furthermore, landscape mosaics enhance disease resilience. Natural buffers act as physical barriers, slowing pathogen spread. Landscape structure determines disease dynamics, with fragmentation disrupting dispersal pathways. Forested corridors filter fungal spores, and diverse non-host plants limit insect vectors. By breaking continuous host plant chains, fragmented landscapes prevent diseases from spreading rapidly, reducing catastrophic crop failure risk.

Thus, viewing natural buffers as integral to agricultural landscapes is crucial for de-risking the supply chain. Well-connected natural habitats in a supply shed provide enhanced pollination, pest control, and disease defense, leading to ecological resilience and economic stability, ensuring a consistent flow of fresh fruit bunches while meeting demand for environmentally responsible production.

Landscape connectivity is a well established approach to quantify landscape diversity and resilience, and could be done using a modelling framework like Omniscape[2] and generate a landscape connectivity score for each and every supply shed.

[2] https://docs.circuitscape.org/Omniscape.jl/latest/

5.2.6 Water Stress - Hydrology Missing from the Picture

While meteorological data related to precipitation, evapotranspiration and soil moisture offers a vital snapshot of climate-driven stress, it is the **hydrological response of the supply shed** that dictates a mill's operational resilience and the viability of its commodity production. Modeling the landscape's water flow – its baseflow, quickflow, and erosion dynamics – transforms a static view of "dryness" into a dynamic blueprint of tangible financial risks. A prolonged drop in **river baseflow,** for instance, is not a theoretical risk; it is a direct threat that forced mills across Sumatra and Kalimantan to halt production during recent El Niño droughts due to critical shortages of process water. Simultaneously, this diminished flow capacity turns a legally compliant effluent discharge into a concentrated pollutant plume, a reality that has led to fines and temporary license suspensions for mills along Malaysia's Kinabatangan River.

Conversely, poor land management that increases **quickflow** turns moderate rainfall into a recurring logistical nightmare. The intense, flashy floods seen in Johor, Malaysia, in early 2024 are a vivid example, washing out the very roads needed to transport fresh fruit bunches (FFB) and directly impacting production targets. A delay of just 24 hours can cause fruit spoilage and a significant drop in the oil extraction rate (OER), directly eroding profitability. This same destructive quickflow drives **soil erosion**, a slow-burning risk quantified by the Malaysian Palm Oil Board as causing yield losses of up to a ton of FFB per hectare annually. This degradation forces a costly reliance on fertilizers to maintain the long-term productivity of the supply shed's core asset. Therefore, modeling a supply shed's hydrology is not an academic exercise; it is a fundamental tool for de-risking operations against these proven, costly shocks.

The best approach for modeling hydrology at the palm oil supply shed level would be to use the **SWAT+ (Soil & Water Assessment Tool)**, a process-based model that creates a "digital twin" of your river basin to simulate key water risks. This approach provides the detailed, dynamic understanding of water availability and soil erosion necessary for a comprehensive supply chain risk assessment. For a faster, less data-intensive overview, the Python-native **InVEST Water Yield model** offers a simpler alternative, ideal for initial strategic comparisons of land-use scenarios.

5.3 The Challenges of Large Scale Validation

5.3.1 Statistical Framework for Validating Environmental Variables at Scale

This study did not formally include a validation strategy due to resource constraints. However, when producing data of this granularity and detail, building confidence in the quality of the data through validation is primordial, without which this data will not make its way into any decision-making process.

To address the challenge of validating environmental variables across vast geographies like the entirety of Indonesia's supply sheds, a robust statistical framework is essential to ensure accuracy while minimizing the operational burden of field data collection. Epoch developed a methodology[3] that employs a model-based, **stratified random sampling** approach that is both statistically powerful and operationally flexible.

The process begins by leveraging "**prior knowledge**" – such as satellite-derived biomass maps or historical land-use change data – to understand the spatial distribution of the variable of interest. This information is fed into a **K-means clustering algorithm** which divides the landscape into distinct zones, or strata, effectively grouping similar areas together. This stratification is crucial as it drastically reduces the number of samples needed compared to non-stratified methods.

Once the strata are defined, the framework uses **Neyman Allocation** to calculate the optimal number of samples per stratum. This method intelligently directs sampling efforts toward strata with higher internal variability, ensuring that resources are not wasted over-sampling large, homogenous areas. A key operational advantage is that the stratification is decoupled from the final sample placement. This allows field teams the flexibility to relocate an inaccessible sample point to an alternative location *within the same stratum*, maintaining statistical integrity while accommodating real-world logistical challenges. This robust and efficient methodology provides a scalable framework for validating a wide range of environmental estimates – from water stress and biodiversity metrics to LUC emissions – across entire supply landscapes.

5.3.2 Supply Shed Extent Validation

This sampling approach is ideal for validating a probabilistic supply shed by treating the model's output as the "prior knowledge" needed to design an effective field survey. The goal is to ground-truth whether the modeled probabilities of sourcing match the reality on the ground.

[3] Epoch, 2025. Optimized Sampling Design. Design for Biomass and Soil-related Field Data Collection.

1. **Stratification Based on Probability:** The probabilistic supply shed map, which shows the likelihood of sourcing from any given area, would be used as the primary input for the **K-means clustering algorithm**. This would create strata based on sourcing probability (e.g., Stratum 1: 80-100% probability, Stratum 2: 60-80% probability, etc.).

2. **Sampling and Data Collection:** After calculating the number of samples per stratum using **Neyman Allocation**, random points representing farms or plantations would be generated. The "field data collection" at these points would not be a physical measurement but a **structured survey**. Field teams would interview the plantation owner or manager at each sampled location and ask the key validation question: "**Which mill do you primarily sell your Fresh Fruit Bunches (FFB) to?**"

3. **Validation:** The survey answers are then compared against the original probability map. If, for example, 9 out of 10 surveyed farmers in the "80-100% probability" stratum confirm they sell to the target mill, it validates the model's high accuracy in that area. If only 5 out of 10 do, it reveals a significant overestimation by the model, providing a quantitative measure of its error.

5.3.3 Environmental Metrics Validation

For Deforestation and LUC emissions data (i.e. ecosystem carbon stock at a given time), this sampling design is a naturally fitting approach. In fact, it was developed for the purpose of calibrating and validating ecosystem carbon stock models for emissions accounting purposes in the context of SBTi and the GHG Protocol Land Sector Removal Guidance (GHGP LSRG).

While the sampling design could be applicable to other biophysical and environmental indicators produced, validating indicators like evapotranspiration anomalies or hydrological model outputs (e.g., from SWAT+), the physical state of water in the landscape has to be measured.

- **Prior Knowledge:** A map of the water stress index would be used to create strata (e.g., "high stress," "medium stress," "low stress").
- **What to Sample in the Field:**
  - **Soil Moisture:** The most direct validation. Using a TDR (Time-Domain Reflectometry) probe or by taking soil cores to measure gravimetric water content at various depths. This directly validates satellite-derived soil moisture products.
  - **Streamflow:** Measuring the discharge rate (volume per second) in local streams and rivers using a current meter. This is essential for validating the outputs of hydrological models like SWAT+, specifically their predictions of baseflow and water yield.
  - **Groundwater Level:** Measuring the depth of the water table in existing wells or by installing piezometers. This provides crucial ground-truth data for validating the groundwater components of a hydrological model.

To validate indicators derived from landscape structure, such as canopy height heterogeneity or connectivity, you must measure the actual presence and diversity of life.

- **Prior Knowledge:** A map of the biodiversity indicator (e.g., a landscape connectivity "current" map from Omniscape) would be stratified into zones like "high connectivity corridor," "moderate connectivity," and "isolated patch."
- **What to Sample in the Field:**
  - **Vegetation Structure:** Establishing sample plots to conduct a forest inventory, measuring tree **Diameter at Breast Height (DBH)**, canopy cover, and understory plant species richness. This directly validates if high "canopy height heterogeneity" truly corresponds to high structural diversity on the ground. This is already sampled in the context of a traditional ecosystem carbon survey, whereby DBH, canopy cover and species inventories are collected.
  - **Indicator Species Surveys:** Conducting standardized point counts or transect surveys for key bio-indicator groups like **birds or butterflies**. The richness and abundance of these species can then be statistically correlated with the connectivity or fragmentation metric.

6 Deliverables

6.1 Web App

We published all the results to a web app: https://supply-shed-viz.epoch-sco2-api.com/, where the results will persist and can be consulted at any time.

Figure 34: Screenshot of the app landing page, allowing to browse all the indicators across all facilities and visualize them on a map or on various histograms and pie charts

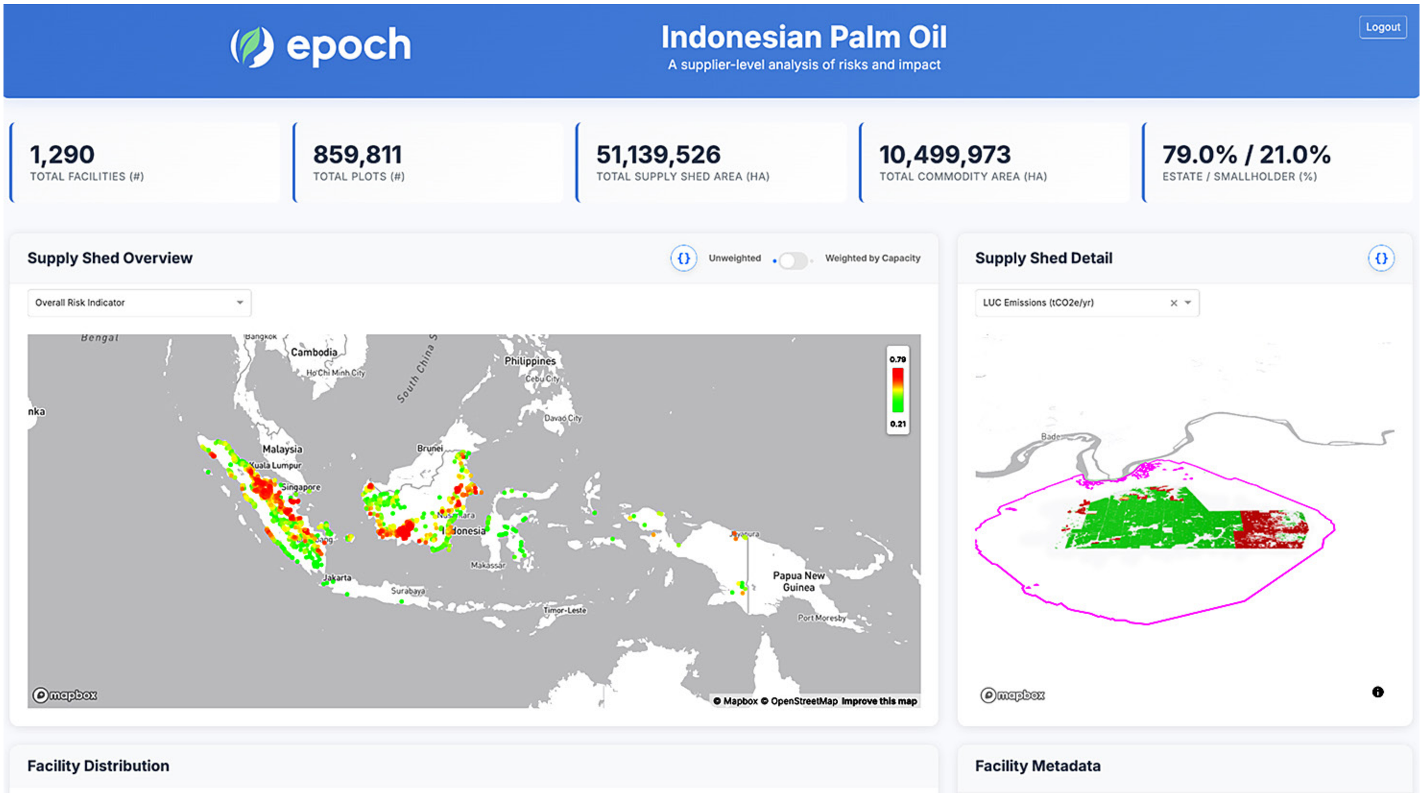

6.2 Files

Geoparquet Files of all the generated supply shed geometries and facilities enhanced with environmental metrics and associated metadata used to produce this analysis have been produced and made available for direct (cloud-native) access at https://storage.googleapis.com/epoch-gee/palm_concessions/supply_shed_geometries.parquet and https://storage.googleapis.com/epoch-gee/palm_concessions/supply_shed_facilities.parquet respectively. This will allow analysts and other more technical users to explore the data in more depth than through the web app interface. If necessary, please email info@epoch.blue for assistance in accessing the files.

7 Conclusion

This pilot study successfully demonstrates a scalable and robust methodology for linking farm-level environmental performance to specific aggregation facilities, fulfilling a core objective of the *Codex Planetarius* initiative. By delineating supply sheds for 1,290 palm oil mills across Indonesia, we have created a direct, traceable connection between over 10 million hectares of commodity production and the first point of aggregation in the supply chain. This powerful framework moves beyond generalized, country-level statistics to provide the granular, actor-specific insights necessary to identify environmental hotspots.

The analysis validates the central premise of *Codex Planetarius:* environmental impacts are not uniformly distributed but are highly concentrated among a small subset of actors. Our findings reveal that the most severe impacts – namely post-2020 deforestation and associated LUC emissions – are overwhelmingly linked to the bottom 5% of facilities, which account for over 73% of the total deforestation detected. This confirms that a targeted approach, focusing on the worst performers, is not only possible but is the most efficient path toward systemic change. By successfully piloting this methodology for Indonesian palm oil, we have established a viable, data-driven foundation for developing the minimum environmental performance standards that *Codex Planetarius* seeks to establish for global commodity trade.

8 Recommendations and Next Steps

The success of this pilot provides a clear pathway to enhance the current analysis, scale the methodology to other critical commodities, and operationalize the *Codex Planetarius* framework. The following steps are recommended to build upon this foundational work.

8.1 Immediate Methodological Enhancements for Palm

While the current indicators provide a strong baseline, their robustness can be significantly improved by incorporating the more advanced modeling techniques identified in the discussion:

- **Refine Emissions Accounting:**
  - **Processing Emissions:** Systematically identify mills with and without methane capture systems (biogas domes) using high-resolution imagery to apply more accurate processing emissions factors.
  - **Historical LUC Emissions:** Extend the LUC emissions analysis back to 2000 to align with the 20-year accounting principles of the GHG Protocol Land Sector and Removals Guidance, providing a more complete picture of legacy emissions.
- **Implement Landscape Connectivity Analysis:** Move beyond the plot-level biodiversity proxy by implementing a landscape connectivity model like Omniscape. This will provide a supply-shed-level score for habitat fragmentation, a more direct measure of pressure on biodiversity and ecosystem resilience.
- **Integrate Hydrological Modeling:** Transition from meteorological water stress indicators to a process-based hydrological model like SWAT+. This will provide a more accurate assessment of operational risks such as mill shutdowns from river drought and supply chain disruptions from flash floods.

8.2 Scaling the Framework to Other Commodities

The four-step methodology developed in this pilot is designed to be commodity-agnostic and can be systematically applied to other high-risk supply chains. A strategic expansion should focus on:

- **Automating Facility Identification and Enrichment:** Leverage the proven potential of satellite imagery embeddings to automate the detection of new processing facilities and reduce reliance on manually curated datasets. Further, the lessons learned from Trase and other field-based efforts should be synthesized to create a systematic approach for enriching facility data with ownership information, which is critical for tying impacts to specific economic actors.

- **Creating a Replicable Playbook:** The process of facility identification, supply shed delineation, commodity mapping, and metric generation can be standardized. For each new commodity (e.g., soy in Brazil, cocoa in West Africa), the key variable to adjust for supply shed generation is the travel-time parameter for the supply shed model, reflecting the unique logistics and perishability of each raw material. For each commodity, a different set of risk factors will be identified, which may require the indicator suite to be tailored to that commodity, but deforestation and emissions will remain relevant regardless.

8.3 Operationalizing for *Codex Planetarius* and Supply Chain Engagement

The ultimate goal is to use this data to drive change. The next steps should focus on making the insights actionable for stakeholders by:

- **Implementing a Formal Validation Strategy:** To build confidence and credibility, a formal validation of the supply shed extents and environmental metrics must be conducted. The proposed stratified random sampling framework offers a cost-effective and statistically robust method to achieve this by using targeted field surveys to ground-truth the model's outputs.

- **Engaging with Identified Actors:** The analysis pinpoints the specific facilities, and by extension the companies, associated with the highest environmental risks. This data provides a powerful, evidence-based foundation for direct engagement with the "bottom 10-20%" of performers to drive targeted interventions, sustainable financing initiatives, and policy enforcement.

- **Developing a Tiered Performance Standard:** The percentile-based ranking of all 1,290 facilities demonstrates that a continuous performance curve exists. This data can be used to define initial thresholds for the "minimum environmental performance standards" envisioned by *Codex Planetarius*, creating a clear, data-driven baseline for international trade.

Sources

Benedict, J. J., Carlson, K. M., Febrian, R., & Heilmayr, R. (2023). *Indonesian palm oil exports and deforestation.* Trase Insights. Link to Source

Descals, A., Gaveau, D. L. A., Wich, S., Szantoi, Z., & Meijaard, E. (2024). *Global mapping of oil palm planting year from 1990 to 2021.* Earth System Science Data Discussions. Link to Source

Guye, V. (2024). *Ivorian Cocoa Cooperatives and their Buyers (IC2B) – database v1.0*. Open Data @ UCLouvain. Link to Source

Hengl, T., Consoli, D., Tian, X., Nauman, T. W., Nussbaum, M., Isik, M. S., ... & Harris, N. (2025). *OpenLandMap-soildb: Global soil information at 30 m spatial resolution for 2000–2022+ based on spatiotemporal Machine Learning*. Earth System Science Data Discussions. Link to Source

Hermantoro, H., Ruswanto, A., Gunawan, S., & Putra, A. P. (2024). *Fresh Fruit Bunch (FFB) Processing Modification to Increase Crude Palm Oil (CPO) Rendement*. IOP Conference Series: Earth and Environmental Science, 1297(1), 012084. Link to Source

IPCC. (2014). *2013 Supplement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories: Wetlands*. Intergovernmental Panel on Climate Change. Link to Source

Judijanto, L. (2025). *Integrating Methane Capture and Biomass Utilisation: A Review of Emission Reduction Technologies in Palm Oil Production*. European Journal of Innovative Studies and Sustainability, 1(5), 1-12. Link to Source

Julyanda, M., Sembiring, M. T., & Absah, Y. (2024). *Strategy to Increase the Yield of Fresh Fruit Bunch (FFB) Harvesting Process with Lean Agriculture Method at PTPN IV Palmco Regional II Marjandi Estate*. Asian Journal of Management Analytics, 3(1), 1-10. Link to Source

Lim, K. H., Lim, S. S., Parish, F., & Suharto, R. (Eds.). (2012). *RSPO Manual on Best Management Practices (BMPs) for Existing Oil Palm Cultivation on Peat*. Roundtable on Sustainable Palm Oil (RSPO). Link to Source

Marshall, E., & Caswell, M. (2011). *Biofuels and Land-Use Change: Estimation Challenges*. USDA Economic Research Service, Amber Waves. Link to Source

Mempel, F., Biddle, H., Garcia, A., Loureiro, B., Soares, J., Ribeiro, V., & zu Ermgassen, E.K.H.J. (2025). *Trase Facilities Map*. Trase. Link to Source

Mundt, C. C. (2002). *Use of multiline cultivars and cultivar mixtures for disease management*. Annual Review of Phytopathology, 40, 381-410. Link to Source

Oguoma, O. N., & Nwachukwu, H. C. (2019). *Application of Queueing Theory to Performance Characterization of a Medium Scale Vegetable Oil-Producing Plant in Owerri*. International Journal of Engineering Research and Advanced Technology (IJERAT), 5(2). Link to Source

Potter, L. (2016). *How can the people's sovereignty be achieved in the oil palm sector? Is the plantation model shifting in favour of smallholders?* In *Land and Development in Indonesia* (pp. 315-342). Link to Source

Safitri, L., Galdos, M. V., Comber, A., & Challinor, A. (2024). *Potential for low-emissions oil palm production in Indonesia: insights from spatiotemporal dynamics*. Environmental Research Letters, 19(5). Link to Source

Saibani, N., Muhamed, A. A., Maliami, M. F., & Ahmad, R. (2015). *Time and motion studies of manual harvesting methods for oil palm fruit bunches: A Malaysian case study.* Jurnal Teknologi (Sciences & Engineering), 74(3). Link to Source

Schuchardt, F., Wulfert, K., Darnoko, D., & Herawan, T. (2008). *Effect of new palm oil mill processes on the EFB and POME utilization*. Journal of Oil Palm Research, 20, 115-126. Link to Source

Silalertruksa, T., Gheewala, S. H., Mungkung, R., & Perret, S. R. (2012). *A Comparative Carbon Footprint Assessment of Palm Oil and Cassava Biofuel Feedstocks in Thailand*. Energy Policy, 43, 306-314. Link to Source

Tscharntke, T., et al. (2010). *The functions of tropical forests in a changing world*. In *Tropical Rainforests and Agroforests under Global Change* (pp. 377-393). Springer. Link to Source

Widyastuti, M. T., et al. (2024). *PEATGRIDS: Mapping thickness and carbon stock of global peatlands via digital soil mapping*. Earth System Science Data Discussions. Link to Source