

Master's thesis and Internship : Energy transition pathways for isolated regions in the Global South - A case study of Bolivia's Northern Amazon

Auteur : Bernaerts, Valentine

Promoteur(s) : Quoilin, Sylvain

Faculté : Faculté des Sciences appliquées

Diplôme : Master : ingénieur civil en génie de l'énergie à finalité spécialisée en Energy Networks

Année académique : 2025-2026

URI/URL : <http://hdl.handle.net/2268.2/27603>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.



LIÈGE université
School of Engineering

University of Liège
Faculty of Applied Sciences
Academic year 2025-2026

Energy transition pathways for isolated regions in the Global South

A case study of Bolivia's Northern Amazon

Master's thesis completed in order to obtain the degree of:

Master of Science in Energy Engineering, professional focus in Networks

Supervisors:

Sylvain Quoilin

Sergio Luis Balderrama Subieta

Co-supervisors:

Claudia Sánchez Solís

Pablo Jimenez Zabalaga

Valentine Bernaerts

Acknowledgements

This thesis marks the end of my engineering studies, and I would like to thank the people who made it possible.

Firstly, I would like to thank my supervisor, Sylvain Quoilin, for trusting me with a subject and then letting me decide how to approach it. He prevented me from spreading my efforts too widely, and our meetings shaped the structure of this document. I am also grateful to him for giving me the opportunity to spend part of this year working in Bolivia, the most interesting experience of my studies.

My thanks also go to Sergio Luis Balderrama Subieta, who welcomed me to the Centro de Investigaciones en Energía at the Universidad Mayor de San Simón, making my four-and-a-half months in Cochabamba possible. He guided the direction of this work and was always available whenever I needed his advice.

I would also like to thank my co-supervisors, Claudia Sánchez Solís and Pablo Jimenez Zabalaga, who taught me the two models that this thesis brings together. Claudia guided me through everything related to demand and the sufficiency framework, while Pablo did the same for EnergyScope MultiCell. They were always available to answer my questions, whether in person, via email or by message, and they kindly took the time to read this manuscript and provide feedback. The thesis would look very different without them.

I am also grateful to everyone who contributed to making my time in Bolivia so enjoyable. On my very first day, Claudia Sánchez Solís and Carlos Fernandez Vazquez picked me up from the airport and helped me with everything that a foreign student cannot manage alone. The whole team at the Centro de Investigaciones en Energía welcomed me on arrival, answered the questionnaire that informed the clustering in Chapter 4, and taught me a lot of Spanish. I am equally thankful to those who took me to see the electrical installations of an isolated mountain community, where I could see for myself the constraints that this thesis describes. That stay was made possible by a grant from the Académie de Recherche et d'Enseignement Supérieur, which I gratefully acknowledge.

My thanks go to Roger Arias Valentines, whose thesis forms the basis of this one, and to Antoine Mouchamps, who read the most important sections and highlighted what was unclear.

Finally, I would like to thank my friends, who supported me over the years, and my family, who have always been there for me. My deepest gratitude goes to my parents, and I know how much I owe them.

The writing of this thesis was assisted by an LLM, Claude, used for drafting, rephrasing and proof-reading in English, and by Claude Code for part of the post-processing and figure scripts. Every modelling choice, every result and every conclusion reported here is my own, and every figure comes from the model runs described in the text.

Abstract

The Bolivian Northern Amazon covers 21 municipalities and 84,209 households. A quarter of them have no public electricity supply, and the region produces almost all of its electricity from diesel. Those households have never been metered, so their demand appears in no record, and no line reaches them, so serving them costs money that a generation model does not see.

This thesis asks what alternative energy transition pathways towards universal access cost and emit for the region. Universal access here means that no household is left without electricity, whether from a line or from its own solar system, and not that every household is connected. RAMP, a stochastic bottom-up model, builds the demand of the unmetered households from the appliances they own. EnergyScope MultiCell, a linear programming model of a whole energy system, finds the cheapest mix of generation, storage and imports that supplies it. A geospatial layer compares, community by community, a medium voltage line with standalone solar home systems, and prices the network that follows. Five scenarios are built for 2025, and seven further runs extend the analysis to 2035 and 2050, one of them with an announced electricity import from Brazil. The comparison separates three levers: how the electricity is produced, the access given to the households that have none, and the level of service assumed for the households that already have electricity.

Rebuilding the supply is the largest lever, and it asks nothing of the population. Replacing diesel generation with utility solar and battery storage cuts the annual cost of the system from 66.51 to 42.15 million euros and reduces its emissions by two thirds. However, waiting is expensive: by 2035 the annual cost of a supply that cannot be rebuilt exceeds that of a rebuilt one by 45.39 million euros, and by 2050 by 70.85 million, because it keeps paying for fuel while the demand grows.

Serving the 20,714 households that are badly served or not served at all costs 6.50 million euros more per year, of which 3.99 million is medium voltage line and 2.51 million new generation. A study that prices only the power plants therefore reports less than half of what access costs.

The level of service assumed for households with electricity is more important than access itself. Bringing every household of the region to the same sufficiency bundle, a normative minimum of energy services, costs 35.27 million euros per year. This is approximately half of what the region currently spends, because the connected households consume more than that bundle. Emissions follow the supply and not the demand: the scenarios that replace the diesel plants emit 58 to 74 kilotonnes of CO₂ equivalent per year, and those that keep them about 217. Rebuilding the supply should therefore come first, and the access programme should be planned alongside it. Only the storage decision should wait for the Brazilian interconnection, since electricity available at every hour does what a battery does.

Contents

Acknowledgements	1
List of abbreviations	1
1 Objective and scope	2
1.1 Context and motivation	2
1.2 Research question and objectives	2
1.3 Thesis structure	3
2 Literature review	4
2.1 Planning rural electrification under data scarcity	4
2.2 Estimating a demand that has never been measured	4
2.3 Optimising the generation mix, from one village to a whole system	5
2.4 Putting the network back into the picture	6
2.5 Contribution of this work	7
3 Methodology	7
3.1 Pipeline overview	7
3.2 Demand modelling with RAMP	9
3.3 EnergyScope MultiCell	11
3.3.1 Model overview	11
3.3.2 Home system technologies	12
3.3.3 Brownfield treatment of existing capacity	12
3.3.4 Water heating in the decentralised heat layer	13
3.4 Constraining production volumes	14
3.4.1 An absolute floor on yearly production	14
3.4.2 An absolute ceiling on yearly production	15
3.4.3 Why a production bound rather than a capacity factor	15
3.4.4 Implementation in the model	16
3.4.5 Scope of the contribution	16
3.5 Spatial electrification layer	16
3.6 Projecting the system to 2035 and 2050	17
3.6.1 Projecting the population and the number of households	17
3.6.2 Projecting who has electricity, and by what means	18
3.6.3 Projecting the demand	19
3.6.4 Costs, inheritance between horizons, and what may be compared	20
4 Case study: the Bolivian Northern Amazon	21
4.1 Regional context and study area	21
4.2 Census 2024 baseline	23
4.2.1 The three populations of the region	24
4.2.2 Infrastructure allocation rules	24
4.3 Scenario definitions	25
4.4 Clustering into EnergyScope cells	26
4.4.1 Clustering variables	26
4.4.2 Weight calibration	27
4.4.3 Number of clusters	27
4.4.4 Clustering results	28
4.5 RAMP input adaptations	29

4.5.1	Thermal comfort: fans calibrated by temperature data	30
4.5.2	Flour processing replaced by rice processing	32
4.5.3	Radio power correction	32
4.6	Sufficiency demand	32
4.7	Current demand	34
4.7.1	Source A: grid-connected demand from utility data	34
4.7.2	Source B: off-grid households modelled with RAMP	35
4.7.3	Source C: non-electrified households	36
4.8	EnergyScope configuration	37
4.8.1	Cluster-independent parameters	37
4.8.2	Inter-cluster exchanges	37
4.8.3	Demand file	38
4.8.4	Time series	40
4.8.5	Resources	41
4.8.6	Technologies	41
4.8.7	Current system configuration	42
4.8.8	Universal access configuration	44
4.9	Spatial layer application	45
4.9.1	Data sources	45
4.9.2	Cost-breakeven classification	46
4.9.3	Source B and Source C integration	48
4.9.4	Standalone technology costs	51
4.10	The 2035 and 2050 horizons	52
4.10.1	Calibrating and testing the projection rules	52
4.10.2	Population and households in 2035 and 2050	54
4.10.3	The split between the three sources at each horizon	56
4.10.4	The projected demand of the connected households	58
4.10.5	The four projection scenarios	60
4.11	What this chapter establishes	61
5	Results	62
5.1	How the costs are reported	63
5.2	Current system	63
5.3	Supply transition	65
5.4	Universal access	67
5.5	Access transition	70
5.6	Universal sufficiency	71
5.7	The hot water technology lever	76
5.8	Comparing the five 2025 scenarios	78
5.9	Results for 2035 and 2050	83
5.9.1	The cost of a locked supply	84
5.9.2	The price of the calendar	85
5.9.3	What the model builds	85
5.9.4	The Brazilian import	86
5.9.5	Sensitivity of the 2050 results	87
6	Discussion	89
6.1	Rebuilding the supply comes first, and the order has a price	89
6.2	The cost and the access are not in the same municipalities	90
6.3	The thesis prices the system, not the bill	90

6.4	The Brazilian import is cheap, and everything else about it is undecided	91
6.5	The cheapest system needs a solar collector on every roof	92
6.6	What would hold in another region	92
7	Conclusion	93
8	Limitations and future work	94
8.1	Future work	95
	Appendices	96
A	Clustering details	96
B	Demand inputs and RAMP profiles	97
C	EnergyScope configuration tables	100
D	Projection details	101
E	Additional figures	103

List of abbreviations

AETN	Autoridad de Fiscalización de Electricidad y Tecnología Nuclear	LNG	Liquefied natural gas
ANH	Agencia Nacional de Hidrocarburos	LPG	Liquefied petroleum gas
BNA	Bolivian Northern Amazon	MHE	Ministerio de Hidrocarburos y Energías
CNDC	Comité Nacional de Despacho de Carga	MTF	Multi-Tier Framework
COP	Coefficient of performance	MV	Medium voltage
CRF	Capital recovery factor	PEVD	Programa de Electricidad para Vivir con Dignidad
CSP	Concentrated solar power	PV	Photovoltaic
ENDE	Empresa Nacional de Electricidad	RAMP	Remote Areas Multi-energy load Profiles
ESMC	EnergyScope MultiCell	RUE	Registro de Unidades Educativas
GIS	Geographic information system	SHS	Solar home system
GWP	Global warming potential	SIN	Sistema Interconectado Nacional
HVAC	High voltage alternating current	TD	Typical days
IDTR	Infraestructura Descentralizada para la Transformación Rural	YPFB	Yacimientos Petrolíferos Fiscales Bolivianos

1 Objective and scope

1.1 Context and motivation

The Bolivian Northern Amazon (BNA) covers the entire department of Pando, the northern municipalities of Beni and the municipality of Ixiamas in La Paz, totalling 21 municipalities. The region holds 317,517 inhabitants and 84,209 households, and a quarter of them remain outside the public electricity service: 11,389 households have no electricity supply at all, and another 9,325 rely on a private solar panel or generator [1]. The detailed demographic baseline is presented in Section 4.2. Most municipalities operate on diesel-based microgrids or have no formal electricity supply at all, and the region burns about 512 GWh of diesel a year to produce its electricity. Road access across the region is limited during the wet season (November to April), when heavy rainfall renders many routes impassable [2].

This situation defines the planning problem addressed in this thesis. The region has to move from diesel towards a reliable and low-carbon supply while reaching the households that have none, and what that costs depends on two choices: the level of service assumed for each part of the population, and whether the existing plants are kept or replaced. Different answers to those two choices lead to different systems, and this thesis prices and compares them.

Data is another constraint. The demand of these households has never been measured, and the equipment they own is known only as a municipal average. Their communities have been mapped, but with data from 2012. Therefore, most of what a planner needs has to be built before anything can be optimised, and this thesis does that work first.

1.2 Research question and objectives

Two features of the region shape the question this thesis asks. A quarter of its households have never been metered, so their demand has to be built before it can be served. None of them is connected to the public service today, so reaching them costs money that a generation model does not count, whether that money buys a line or a standalone system. Both features are common across the Global South, and neither is represented in the energy system models this thesis uses, which treat the distribution network as already built and the demand as measured. The main research question is:

What are the costs and the emissions of alternative energy transition pathways towards universal access in the Bolivian Northern Amazon?

Two terms of that question carry a precise meaning here. Universal access means that no household is left without electricity, whether it receives that electricity from a line or from its own solar system. It does not mean that every household is connected to a grid. A sufficiency bundle is a normative list of the energy services a household needs for a decent standard of living, taken from the literature and used here as one of the two demand definitions that are compared.

The question is evaluative, not prescriptive. The thesis compares the pathways on their annual cost and on their emissions, and Chapter 7 closes on recommendations that follow from that comparison.

This thesis chains two models and a spatial post-processing layer to answer that question. RAMP, a stochastic bottom-up demand model, builds the demand of households that have never been metered. EnergyScope MultiCell, a linear programming model of the whole energy system, finds the cheapest generation mix for that demand. The spatial layer adds the cost of reaching them, whether by extending a line or by installing a standalone system. The region is represented through five clusters, and

that resolution keeps the main spatial differences while leaving the optimisation tractable.

Arias built this chain for Bolivia as a whole [3]. This work applies it to one region and keeps its municipalities apart instead of merging them, and that resolution makes the problem visible. The share of households connected to the public grid runs from 0.3 to 95.9 % between two municipalities of the same region, and a regional average would erase the very problem this thesis studies.

The question is addressed through three objectives:

Objective 1. Demand characterisation.

Estimate the electricity demand of the region under two definitions: what the population consumes today, reconstructed separately for the three groups that the census distinguishes, the households connected to the public service, those supplied by their own panel or generator and those with no supply at all, and what a normative sufficiency bundle would require of every household.

Objective 2. Generation mix optimisation.

Determine, for each combination of demand and investment rules, the cost-optimal technology mix, storage and external supply, and identify which components of the annual system cost respond to the demand definition and which respond to the supply strategy.

Objective 3. Transition pathway.

Establish the full cost of universal access, including the lines or the standalone systems needed to reach the households that have none. Then follow the region to 2035 and 2050: compare a supply that stays on diesel with one that is rebuilt, compare two dates for reaching every household, and test what an electricity import from Brazil changes.

1.3 Thesis structure

The document follows the order of the work itself. Chapter 2 reviews how rural electrification is planned where measured data is scarce, how a demand that has never been measured is built, and how generation mix optimisation and spatial electrification tools split the planning problem, one resolving the supply and the other the space. It closes by stating what this thesis adds to each of those three strands.

Chapter 3 describes the two models and the spatial post-processing layer that are chained here: RAMP for the demand, EnergyScope MultiCell for the generation mix, and the spatial layer for the network. Section 3.4 sets out the change made to the optimisation model, without which the access question could not be expressed. Section 3.6 states the rules that carry the whole chain to a future horizon.

Chapter 4 applies all three to the region. It builds the census baseline and the three populations used throughout, defines the five scenarios, groups the 21 municipalities into five clusters, adapts the demand inputs to the Amazon, and configures the model cluster by cluster. Section 4.9 classifies the 697 communities of the region between a grid extension and a standalone system, and Section 4.10 carries the region to 2035 and 2050.

Chapter 5 prices the five scenarios of 2025 and decomposes the distance between the present situation and universal sufficiency into three steps, then reports the seven runs of 2035 and 2050 and compares two access calendars against a region that keeps the supply it has today.

Chapter 6 reads these results from the point of view of someone who has to decide what to build, and states what the model leaves out. Chapter 7 gathers the answer to the research question and closes on three recommendations. Chapter 8 gathers the limitations that weigh most on the results, the others being stated where they arise, and closes on the work this thesis makes possible.

2 Literature review

2.1 Planning rural electrification under data scarcity

Around 666 million people in the world still had no access to electricity in 2023, and the pace of progress is not enough to close that gap by 2030 [4]. What remains of the problem is also the hardest part of it. The households still unconnected live in remote, low income and often fragile areas, which is exactly where the cost per connection is highest and where the data needed to plan is thinnest.

Planning for those households means choosing between three options for every settlement: extend the existing network, build a local mini-grid, or install standalone systems. The choice is economic and it depends on distance, on the number of households and on how much electricity they will use. None of these three quantities is easy to obtain in the places where the question arises.

The binding constraint is the data rather than the method. A household that has never been connected has no consumption record, so there is nothing to fit a demand model on. Utility statistics cover the systems that exist and are silent on everything else, which makes the unelectrified population statistically invisible. In the municipal consumption database of [5], five municipalities of the study area report zero measured consumption, not because nobody consumes anything there, but because private panels and generators are not metered. Access is also a matter of degree. The Multi-Tier Framework formalises this by grading it against the power a household can actually draw [6].

Bolivia is a good illustration of the gap between a national average and a local situation. The country has raised its coverage substantially over two decades, and the remaining deficit is concentrated in the rural lowlands and in the Amazon. The BNA is the extreme case: a quarter of its households are either unconnected or supplied by their own equipment, and the coverage of its municipalities spans almost the whole scale, as Figure 4 shows.

Most published work sits at one of two scales. National models cover the whole territory but aggregate the rural population into large zones, which erases that heterogeneity [7, 3]. Community scale studies keep the detail but describe one village, which is too small to inform an investment programme [8, 9]. The intermediate scale, a region large enough to be planned and small enough to keep municipalities distinct, is where this thesis works.

2.2 Estimating a demand that has never been measured

Where consumption records do not exist, demand has to be built from the appliances people own and the hours they use them. This demand is latent: it exists, it cannot be measured, and households satisfy it as far as their own means allow. Mandelli et al. formalised this as a stochastic bottom-up procedure and implemented it in LoadProGen, drawing each appliance switch-on at random within a declared time window so that the aggregate profile of a community keeps the diversity of a real load [10]. RAMP extends the same idea, adds randomness on power levels and duty cycles, and covers services that are not purely electric such as water heating and cooking [11, 12]. Section 3.2 describes the version used here.

RAMP was validated against measurements taken on the El Espino microgrid in Bolivia, which is what makes it a defensible choice for this region rather than a generic one [11]. It has since been calibrated for the main rural zones of the country, with appliance sets and usage windows built from field surveys and national reports [13], and the lowlands configuration of that work is the starting point of the inputs adapted in Section 4.5.

Building the profile is only half of the problem. The other half is deciding what level of service to put into it, and the literature answers that question in two incompatible ways.

The first answer is normative. It states what a household should receive and models that, whatever it consumes today. Sanchez-Solis et al. define energy sufficiency as the set of services needed for a decent standard of living, and derive electricity thresholds of 393 to 653 kWh per person and year across five rural cases of the Global South, with cold storage as the dominant contributor and thermal comfort as the main source of variation between climates [14]. Similar thresholds exist elsewhere, expressed as a minimum annual consumption rather than as a list of services [15]. The appeal of this approach is that it is a target. Its weakness is that it says nothing about the distance between the target and the present situation.

The second answer is descriptive. It reproduces what is consumed today in regions with similar conditions but with access to any form of supply, and projects it forward, which is what any model calibrated on utility data does. It is grounded in measurement, and it inherits the inequality of the situation it measures. In a region as unequal inside itself as this one, a demand built on the trend keeps that gap open.

The two are rarely put side by side. Studies that soft-link a demand model to a generation mix optimisation vary the demand along socio-economic trajectories rather than along the definition itself [16]. Work in the sufficiency line builds the normative demand and optimises against it, without a measured counterfactual on the same territory [14, 3]. The comparison matters because the two definitions do not only differ in level. A normative bundle is uniform across households and a measured demand is not, and Section 5.8 shows that this changes which lever reduces the system cost the most.

One population falls between the two, and neither of the two lines of work represents it explicitly. Households with a solar panel or a small generator are neither unelectrified nor properly served, and no census records what they can actually run. The treatment adopted here combines the appliance ownership of the census with the power ceilings of the Multi-Tier Framework [6], so that a solar home system is not credited with a refrigerator it cannot start. Section 4.7.2 gives the resulting retention factors.

2.3 Optimising the generation mix, from one village to a whole system

Once a demand exists, a model has to decide what satisfies it. Two families of tools do this and they were built for different questions.

The first sizes a single system. Given a load profile and a solar resource, it finds the combination of panels, batteries and generators that supplies that load at the lowest net present cost. This is what microgrid design tools do, and they do it with a level of technical detail that a larger model cannot carry. Their limit is the number of sites. Designing every community of a country one by one is computationally out of reach, which is the problem Balderrama addresses by replacing the optimisation of each site with a regression trained on a database of already optimised ones [17, 18]. That work computes a least-cost path to full coverage for the mid-size communities of the Bolivian lowlands, which is the closest published precedent to the question asked here.

The second family optimises a whole system at once. These models cover every sector and every energy carrier over a full year, and they answer where to invest rather than how to build one plant [19]. OSeMOSYS is the most widely used of them [20]. EnergyScope belongs to the same family but describes the system by its end uses rather than by its energy carriers [21]. A household does not demand electricity, it demands light, cold storage and hot water, and the optimiser chooses which

carrier delivers each of them.

That distinction decides the answer in this region. Cooking and domestic hot water are 74 percent of the useful energy of the sufficiency bundle, as Table 8 shows, and neither has to be electric. A model that starts from electricity demand cannot see that choice at all. Section 5.7 measures what it is worth here.

EnergyScope MultiCell extends the formulation to several interconnected cells with their own demand, resources and technologies [22]. The Bolivian branch of this line is well developed. A national model covers the whole energy system to 2050 [7], and a fork adds home system technologies so that dispersed households can be represented inside a national run [3]. Recent work links the same national model to a spatially detailed power system model [23], which shows that the spatial resolution of the Bolivian branch is being improved, but on the transmission side rather than on the last mile.

Two things are missing from this family for the question of this thesis. The distribution network is treated as already built, so reaching a household that has no line today costs nothing in the optimum. The constraints available to hold part of a system fixed are also written as shares of a quantity the optimiser itself decides, which makes them unusable when one population must stay exactly as it is while another is served. Section 3.4 deals with the second, and the next subsection with the first.

2.4 Putting the network back into the picture

The cost of reaching a household is a geographic quantity. It depends on the distance to the nearest line, on how many households share that line, and on how much each of them will draw. Geospatial electrification tools answer this directly. OnSSET compares grid extension, mini-grids and standalone systems for every populated point of a territory and returns a spatially explicit least-cost technology mix, which removes the crude split between connected and unconnected that earlier planning relied on [24].

The approach has been applied to Bolivia. Peña Balderrama et al. generate demand profiles with RAMP, optimise mini-grids, and feed the resulting costs into OnSSET to classify communities [25]. The line and connection unit costs used in Section 4.9.2 come from that work. On the Bolivian Highlands, Stevanato et al. formulate the same trade-off as a mixed integer problem and show that adding the development opportunities lost under a limited service changes the optimal number of connected households by up to a third [26].

Others have connected such a tool to a long-term planning model. Kenya and Ethiopia have both been studied by soft-linking OnSSET to OSeMOSYS, so that the connections implied by an electrification pathway appear in the capacity expansion of the power sector [27, 28]. This is the architecture that comes closest to the one used here.

What these tools represent poorly is the supply. A geospatial model needs a cost per kilowatt hour for each option, and it obtains it from a levelised cost or from a surrogate fitted beforehand. It does not carry a full energy system, so cooking, hot water and the substitutions between carriers are outside its scope. The two families are therefore complementary rather than competing: one resolves space and simplifies supply, the other resolves supply and ignores space.

The layer built in this thesis sits on the simple side of that comparison. It builds the network as a shared tree rather than optimising its routing, and it does not consider intermediate configurations such as a village mini-grid. Its purpose is to supply the two quantities that the optimisation cannot compute on its own, and Chapter 8 states what it leaves out.

2.5 Contribution of this work

The previous subsections leave three openings, and this thesis works in each of them.

The sufficiency literature defines what a household should receive and optimises a system against it [14, 3]. It does not say how far the present situation is from that target, because it carries no measured counterfactual. This work builds both on the same region and with the same model. The current system is costed from utility sales and existing plants, the sufficiency target from the normative bundle, and three intermediate configurations separate the two so that the distance between them can be walked in steps that add up exactly. Section 5.8 gives that decomposition. It shows that the demand definition is not a second order assumption: bringing the connected households down to the bundle saves more than three times what serving the unconnected ones costs.

Expressing that comparison required a change to the model rather than to its inputs. The constraints available in EnergyScope MultiCell bound a technology as a share of a quantity the optimiser itself decides, which fails as soon as one population must stay exactly as it is while another is served. Two absolute bounds on annual production were added, a floor and a ceiling. The floor makes the cost of access measurable against a frozen system, and the ceiling stops a household generator from supplying consumers it cannot physically reach. Section 3.4 states the formulation and why a capacity factor could not play the same role.

Even with those bounds, a whole energy system model assumes the distribution network is already built, so connecting a household that has no line today costs nothing in its optimum. The spatial layer of Section 4.9 removes that assumption. It extends the existing medium voltage network community by community and keeps a line only where that line beats a standalone system, so it classifies the 697 communities of the region and prices the network at the same time. In this region, of the 6.50 M€/yr that universal access costs above the current system, 3.99 are medium voltage lines and only 2.51 are generation. Therefore, a planning study that stops at the generation optimum reports less than half of the real figure. This is a lower bound, since the increment does not cover the replacement of the equipment of the households that stay dispersed.

Three further choices follow from the region rather than from the literature. The analysis works at municipal resolution, because a regional average would hide how far apart the municipalities of this region are. It represents the households that own a panel or a generator as a population of their own, neither unserved nor properly served, by crossing the appliance ownership of the census with the power ceilings of the Multi-Tier Framework [6]. It also keeps cooking and domestic hot water inside the system boundary, and that matters more than expected: the choice between a solar collector and an electric resistance is worth four fifths of the entire medium voltage network in the sufficiency scenario. Section 5.7 quantifies it.

What is new here is the comparison that the chain makes possible on one territory. Each of the three links has a precedent taken on its own [3, 25, 27, 28], and none of those precedents puts a measured demand, a normative one and the network that serves them on the same region and inside the same model.

3 Methodology

3.1 Pipeline overview

Planning the energy transition of an isolated region raises two problems that a standard energy system model cannot solve on its own. The first is the demand. There is no measured consumption data for

households that have never been connected to a grid, so there is nothing to optimise against. The second is the network. Once a least-cost generation mix is found, the model assumes that every household can already be reached, so the cost of building the lines to connect them is simply missing from the result. The methodology of this thesis chains two models and a spatial post-processing layer to deal with both problems. Figure 1 gives an overview.

RAMP generates the demand. It is a bottom-up stochastic model: instead of starting from consumption records, it starts from a list of appliances. Each consumer type, such as a household, a school or a shop, is described by its appliances with their rated power, daily usage time and the time windows in which they are used. RAMP then draws load profiles at one minute resolution, adding randomness to the switching times so that no two days are identical. This makes it suitable for regions where the only available information is what people own and how they use it. The output is one demand profile per consumer type, which is scaled by the number of consumers to obtain the demand of a full community.

EnergyScope MultiCell optimises the generation mix. It takes the annual demand and its hourly profile as inputs, together with the resources and technologies available in each part of the region, and finds the combination of generation, storage and imports that meets the demand at the lowest total annual cost. The region is divided into a small number of cells that can exchange electricity, which keeps the main spatial differences without making the problem too large to solve. The outputs are the installed capacity of each technology, its hourly operation, the total annual cost and the emissions. One thing is missing from this optimum: EnergyScope treats the distribution grid as already built, so connecting a household that has no line today costs nothing in the model.

The spatial layer fills that gap. Using the location of each community and the geometry of the existing medium-voltage network, it compares two options for every community: extending a line from the nearest grid, or installing standalone home systems. The cheaper option per household wins. This classification produces the two numbers that EnergyScope cannot compute on its own. The share of demand that should be served by home systems is fed back into the optimisation, and the cost of the line extensions is added to the generation cost to obtain the full cost of universal access.

The next subsections describe each model in turn, and Chapter 4 applies the full chain to the BNA.

The models pass data through files, so every transfer between two stages is checked against a total computed independently on the other side, and any discrepancy is treated as a defect until explained.




Figure 1: The pipeline of this work. Solid arrows carry the demand and the results, dashed arrows the two quantities the spatial layer feeds back.

3.2 Demand modelling with RAMP

RAMP (Remote Areas Multi-energy load Profiles) is an open source stochastic model that builds an electricity demand from the bottom up [12, 11]. It was designed for the situation this thesis faces: places where consumption has never been measured, and where the only thing that can be described is what people own and roughly when they use it. Rather than extrapolating from records that do not exist, it starts from an appliance list and simulates the behaviour of the people using it. The model has been calibrated and validated for rural Bolivian communities [13].

Each type of consumer is described as a list of appliances. For every appliance the model needs the rated power P , the number of units present, the total time it runs on an average day t , the time windows during which it may run, and the minimum duration of one switch-on c . Two further parameters set how much these can vary: one shifts the boundaries of the time windows from day to day, the other stretches or shrinks the daily usage time by a fraction r_t . From this description RAMP draws a switch-on schedule for every appliance on every day of the year, and sums them into a load profile at one minute resolution.

That schedule is drawn in three steps, and an appliance may also stay off for the whole day. An occasional appliance carries a probability ϕ of being used at all, drawn once per unit and per day, and this is how an ownership rate enters the model. When the appliance is used, its realised daily functioning time is drawn around the declared one:

$$v \sim \mathcal{U}(1 - r_t, 1 + r_t), \quad \tilde{T} = \max(c, \text{round } \mathcal{U}(t, \lfloor tv \rfloor)) \quad (1)$$

The draw has two stages rather than one, since the width of the second interval is itself random. The result is centred on t but concentrated around it, and it is floored at the minimum duration of one switch-on. It is also capped by the total width of the usage windows, which binds only for appliances declared to run almost all day.

The switch-on events are then drawn one after the other. Let F be the set of intervals still free inside the windows, each written $[\alpha, \omega)$. An event may start at any minute that leaves at least c free minutes

in front of it, so the admissible set of start times is

$$A(F) = \bigcup_{[\alpha, \omega) \in F, \omega - \alpha \geq c} \{\alpha, \alpha + 1, \dots, \omega - c\} \quad (2)$$

and the start time and the duration of one event are

$$\tau \sim \mathcal{U}(A(F)), \quad D = \lfloor \mathcal{U}(c, \min(\tilde{T}, \omega_\tau - \tau)) \rfloor \quad (3)$$

where ω_τ closes the free interval that contains τ . The occupied minutes leave F , so two events of the same appliance never overlap, and the draw repeats until the occupied minutes reach $\tilde{T}$. No hour is preferred inside a window: the whole shape of the profile comes from the windows, from the durations and from the aggregation, and never from an imposed time preference.

A power is finally assigned to the occupied minutes. An appliance without a duty cycle carries one level over the whole event:

$$p(m) = \xi n_{on} P, \quad \xi \sim \mathcal{U}(1 - r_p, 1 + r_p), \quad m \in \{\tau, \dots, \tau + D - 1\} \quad (4)$$

where r_p is the variability on the power and n_{on} the number of units switching on together. An appliance with a duty cycle C repeats that cycle over the event instead:

$$p(\tau + j) = n_{on} C [j \bmod |C|], \quad j = 0, \dots, D - 1 \quad (5)$$

so the load of such an appliance is carried by the levels declared in its cycle rather than by the rated power. The refrigerator of this work is in that case, with a compressor level and a standby level alternating inside every switch-on.

The randomness matters more than it might seem. To simulate a community of five hundred households, RAMP does not compute one household and multiply by five hundred. Every consumer is drawn on its own, so the profile of a community is

$$P(t) = \sum_{u \in U} \sum_{i=1}^{N_u} p_{u,i}(t) \quad (6)$$

where U is the set of consumer types, N_u the number of consumers of type u present in the community, and $p_{u,i}(t)$ the profile drawn for the i -th of them. The individual peaks therefore do not line up. What comes out has the smoothing and the diversity of a real community load, with a realistic ratio between the peak and the average. This is what makes the output usable by a system optimisation, since an artificially peaked demand sizes the supply far too large.

Demand is described in three groups: households, community services such as schools, health centres and public lighting, and income generating activities such as shops, restaurants and small workshops. Keeping them separate matters because they load the system at different hours of the day, and because they scale with different drivers. Household demand follows the population, while the other two follow the level of service and economic activity a community has reached.

The last step of the chain is a change of units. RAMP produces electricity, which is what goes into the appliance. EnergyScope works with useful energy, which is the service that comes out of it: the light in the room, the cold inside the fridge, the heat in the water. The two are linked by the efficiency of the appliance, and EnergyScope carries those efficiencies in its own database. Each RAMP category is therefore mapped to an EnergyScope end use, and the annual useful demand of that end use is

$$D_{eut}(c) = \eta_{eut} \sum_{a \in A_{eut}} E_a(c) \quad (7)$$

where A_{eut} is the set of RAMP categories mapped to the end use eut , $E_a(c)$ the annual electricity that category a draws in cluster c , and η_{eut} the conversion efficiency of the appliance that serves the end use. Doing it this way is what leaves the optimiser free to choose which technology, and which energy carrier, delivers the service. The conversion is applied when the demand file is written, so EnergyScope receives useful energy and applies no efficiency a second time. Section 4.8.3 gives the mapping used here.

Two separate inputs are extracted from the same RAMP output. The annual useful demand of each end use sets the level, and a normalised hourly profile sets the shape, telling the model how that annual total is distributed across the 8,760 hours of the year. Splitting level from shape means the two can be adjusted independently, which is what makes it possible to build a scenario where the amount of demand changes but its timing does not.

3.3 EnergyScope MultiCell

EnergyScope MultiCell (ESMC) is an open source linear programming framework that optimises at the same time the investment in an energy system and its hourly operation over a full year [22, 29]. The system is divided into cells. Each one has its own demand, its own resources and its own set of available technologies, and clusters can exchange energy carriers subject to a transfer capacity and a transport cost. In this work a cell is a cluster of municipalities, and the text says cluster from here on. The model code calls the same object a region. The optimiser minimises the total annual cost of the whole system, which is the sum of the annualised investment, the maintenance and the operating cost of the resources used.

Building a model requires three groups of input files: parameters shared by all cells, such as the discount rate, the network losses and the resource prices; exchange parameters describing the distances and transfer capacities between cells; and one set of files per cell describing its demand, resources and technologies. Chapter 4 gives the values used for the BNA.

The version used here is the fork developed for isolated Bolivian communities [3], which adds home system technologies to the standard formulation. The base technology and resource parameters come from the national Bolivian model [30]. The subsections below describe the mechanisms that matter for this work: the home systems, the treatment of already existing capacity, and the decentralised heat layer that serves domestic hot water.

3.3.1 Model overview

The temporal resolution uses 16 typical days selected by K-medoids on the hourly profiles, following the EnergyScope TD methodology [21]. The number of days is kept from the configuration of Arias [3], so the results of the two studies stay comparable. This reduces the 8,760 hours of the year to a tractable set while preserving the variability of demand and renewable supply. The model is implemented in AMPL/Julia.

The model resolves an energy balance. At every hour of every typical day, the supply of each energy carrier must equal its consumption, and the optimiser adjusts the capacities so that the balance is maintained over the whole year. However, this balance is an energy balance and nothing else. The model does not consider power flow, voltage, frequency or reserve requirements, so a mix that balances energy can have operating problems that the model cannot anticipate. This thesis therefore assumes that dispatch is sufficient to cover demand and the flexibility required. Dedicated power system tools can verify that assumption, and a soft-link between the national Bolivian model and such a tool already exists [23]. Applying this to the mixes in Chapter 5 is a natural extension of this work. The year is also represented by 16 typical days, and daily storage returns to its starting level at the end of each of them, so a sequence of several overcast days does not exist in the model and the batteries are sized for the worst typical day rather than for the worst week. Both simplifications point the same way: the systems that rebuild the supply come out smaller and cheaper than a system built to a reliability standard would.

3.3.2 Home system technologies

In the standard formulation all consumers of a cluster draw from a single shared electricity layer, which implicitly assumes that a distribution grid reaches every one of them. In an isolated region this is not true, and a household far from any line has to generate its own electricity on site.

The fork used here adds a dedicated set of home system technologies for this purpose: an individual photovoltaic panel, a small household diesel generator and a dedicated battery. Two things separate them from their utility scale equivalents. They carry no network losses, since production and consumption happen at the same point. And a coupling constraint forces the home battery to be charged only by the home panel, which stops the optimiser from treating it as cheap storage for the rest of the cluster.

How much of the demand of a cluster is served this way is controlled by two mechanisms rather than one, and they work in opposite directions.

The first is a minimum share. The parameter `share_dispersion` requires the home system technologies to cover at least a given fraction of the electricity supplied to the cluster, and the requirement holds for every hour of every typical day. Left at zero it constrains nothing and the optimiser decides freely. Set to a positive value it is a floor, and a floor can only push the production of the home systems up.

The second is a pair of bounds on the yearly production of a technology, written as a quantity of energy rather than as a share. The floor plays the same role as the minimum share but with a fixed right hand side, which the share does not have. The ceiling has no equivalent in the standard formulation, and it is the only way to push the production of a home system technology down rather than up.

The second mechanism answers a problem the first cannot address: a technology belonging to a single household can end up supplying consumers it cannot physically reach. Section 3.4 describes both bounds, and Section 3.5 explains where the values come from.

3.3.3 Brownfield treatment of existing capacity

A snapshot model of an existing system has to represent power plants that are already built. These are given a forced minimum capacity, but charging them the full investment cost of a new plant would be wrong, since that money has already been spent. The investment cost constraint is therefore modified

so that capital expenditure applies only to the capacity built beyond what exists:

$$C_{\text{inv}} = c_{\text{inv}} \cdot (F - f_{\text{min}}) \quad (8)$$

where F is the total installed capacity of the technology and f_{min} the capacity already in place. Existing plants still carry their maintenance cost and still consume fuel when they run, but they add no capital cost to the system. Without this correction the model would charge greenfield capital expenditure on the whole fleet, and the cost of the current system would be strongly overstated. The same rule applies to the home system technologies, where the panels and generators that households already own are credited in the same way.

3.3.4 Water heating in the decentralised heat layer

Domestic hot water is not served on the electricity layer. It is served on a separate low temperature heat layer, where three technologies can meet the demand: a direct electric resistance heater, a solar thermal collector and an electric heat pump. Each of them is coupled to a hot water storage tank, which lets the model produce heat at a different moment from when it is used. The solar collector has no storage of its own, so it is attached to a host technology and shares its tank.

The two layers are linked, since the electric options draw power from the electricity layer to produce heat. The choice of water heating technology therefore changes the electricity demand that the supply side has to cover, and a technology that produces heat directly from the sun removes that demand entirely. This is a power-to-heat coupling between the electricity and heat sectors. Full sector coupling, which would also include transport and industry, is outside the scope of this work.

Figure 2 shows how the pieces fit together for a single household.




Figure 2: Household domestic hot water system. The optimiser chooses how much of the yearly demand each path serves.

One limitation follows from the temporal structure of the model. Hot water demand carries no time series of its own, so the model divides the annual total by the 8,760 hours of the year and holds that value at every hour, in every typical day and every season. The morning and evening draw peaks that RAMP produces are therefore invisible to the heat layer, and the storage they would call for is handled implicitly by the daily tank rather than resolved hour by hour. The same convention applies in every scenario, so the comparisons of Chapter 5 hold. What it affects is the sizing of the tanks, which comes out optimistic, and probably the comparison of Section 5.7, since a flat draw favours the solar collector against the electric resistance.

3.4 Constraining production volumes

Two situations in this thesis require part of the system to be held fixed while the rest is optimised. The first is a question about access. It asks what it costs to supply the households that have no electricity, while the households that already have it keep being served exactly as they are today. The difference in cost against the current situation is then the cost of the access, and nothing else. The second is a question of physical reach. A technology that belongs to a single household must not end up supplying the rest of the cluster. Both require the model to be told how much energy a given technology delivers over the year. EnergyScope MultiCell had no way to say that.

What it had was a minimum share. Before writing it, two quantities are useful. The yearly production of a technology j in a cluster c is its hourly output summed over the typical days, each hour weighted by the number of real hours it stands for:

$$P(c, j) = \sum_{t \in \text{PERIODS}} F_t(c, j, h_t, td_t) t_{op}(h_t, td_t) \quad (9)$$

The yearly supply of an end-use layer eut is the same sum over every technology that serves it, plus what the resources bring to that layer:

$$S(c, eut) = \sum_{j' \in T_{eut}} P(c, j') + \sum_{r \in \text{RES}} \sum_t \ell_{r, eut} \left(R_t^{loc} + R_t^{ext} + R_t^{imp} - R_t^{exp} \right) \quad (10)$$

where $\ell_{r, eut}$ is the conversion coefficient of resource r into layer eut , and the four R terms are the local, external, imported and exported use of that resource. The existing constraint then forces a technology to cover at least a given fraction of that supply:

$$P(c, j) \geq f_{min}^{perc}(c, j) \cdot S(c, eut) \quad (11)$$

The problem is that j is itself part of S . Forcing the technology to produce more raises S , which raises the floor it has to meet. The target moves as soon as it is pushed.

This showed up clearly in practice. The production always landed below the intended value, and the shortfall was the same percentage for every technology sharing a layer. A wrong input would not do that. A shared right hand side would. Recalibrating the percentage from the observed result closed part of the gap at each attempt, but never all of it. The limitation is in the shape of the constraint, not in its calibration.

3.4.1 An absolute floor on yearly production

The model was therefore given a second option. A parameter $f_{min}^{prod}(c, j)$ sets a minimum yearly production in GWh, and the constraint becomes:

$$P(c, j) \geq f_{min}^{prod}(c, j) \quad (12)$$

The right hand side is now a number. It does not depend on anything the optimiser decides, so the target is met exactly. The constraint fixes how much a technology produces over the year. It does not

fix when. The optimiser can still move that production between hours, and it stays free on the rest of the system.

3.4.2 An absolute ceiling on yearly production

A floor alone is not enough. The same mechanism is needed in the other direction, for a reason that comes from the structure of the model rather than from the data.

A cluster has a single electricity layer. Every consumer of the cluster draws from it, and the model has no way to tell apart a household connected to the grid from one that is not. The home system technologies of Section 3.3.2 are correctly excluded from the network losses, because production and consumption happen at the same point. That exclusion has a side effect. Their delivered kilowatt hour is cheaper than the one delivered by a utility scale plant of the same fuel, which pays the losses. The optimiser sees a cheap supplier and uses it. It runs the generators of off-grid households to serve demand elsewhere in the cluster.

This is physically impossible. A household that is not connected to any line can only supply itself. What the shared layer cannot represent is a topology, and no price signal will restore it.

The parameter $f_{max}^{prod}(c, j)$ closes this, symmetrically to the floor:

$$P(c, j) \leq f_{max}^{prod}(c, j) \quad (13)$$

Setting the ceiling equal to the demand of the households that own the technology restores the balance that the missing topology would have imposed. The value is not fitted to a target cost. It follows from the definition of the population itself, since a household counted as off-grid is one that no line reaches. The correction therefore states a fact about the system.

The bound constrains a volume, not a destination. On a shared layer the model still cannot say which household received which kilowatt hour. What it can guarantee is that the technologies owned by off-grid households deliver over the year exactly the energy those households consume. That remaining ambiguity is a property of the single layer. Every statement in this thesis about home systems is therefore a statement about volumes and not about topology. Representing the off-grid population properly requires a second electricity layer per cluster, which is a change to the shared model code rather than to its inputs.

3.4.3 Why a production bound rather than a capacity factor

The model already carries a parameter that limits the yearly production of a technology. It is called c_p , and the constraint it enters bounds exactly the quantity of Equation 9:

$$P(c, j) \leq F(c, j) \cdot c_p(c, j) \cdot T \quad (14)$$

where $F(c, j)$ is the installed capacity of the technology and T the total number of hours the year represents. Lowering c_p so that the technology simply runs less was the obvious route. It was not taken, and the reason is not the same depending on whether the capacity is fixed or free.

Where the capacity is locked to what already exists, F is a constant and the right hand side is a number. The bound would work. But obtaining a chosen production would mean solving backwards for c_p , which is setting a technical parameter from the answer it is supposed to produce. The value

would describe nothing about the machine and would not transfer to any other study. There is nothing to calibrate it against either, since the measurements available in the literature describe backup units or community generators rather than a domestic primary supply [9].

Where the capacity is free, the problem is a different one, and it is the problem already met above. F is a decision variable, so the right hand side moves with whatever the optimiser chooses to install. Pushing the production up raises the capacity, which raises the bound. This is the moving target of Equation 11 in another form.

That is the common thread between the two routes that were not taken. Neither of them constrains a production, because neither has a constant right hand side. One is a fraction of the supply of a layer, the other a fraction of an installed capacity, and the optimiser decides both. The bounds added here are absolute.

3.4.4 Implementation in the model

Both bounds are inactive by default, the floor at zero and the ceiling at a value far above any production the system could reach. A scenario that sets no value is therefore left untouched.

The constraint rows are also generated only for the technologies that carry an explicit bound. A technology left at its default produces no row at all, rather than a row whose right hand side lies many orders of magnitude outside the range of any feasible production. This matters because the model already spans a very wide range of magnitudes, from household scale capacities of a few hundred watts to regional energy flows of hundreds of gigawatt hours. Keeping the numerical range of the problem close to the range of its meaningful values costs nothing to implement and avoids adding solver noise that would be invisible in the results.

3.4.5 Scope of the contribution

The difference between constraining a share and constraining a volume does not matter when the whole system is optimised at once, which is the usual setup in a study of the future. It matters as soon as one part of the system has to stay untouched, or as soon as a technology has to be prevented from serving a demand it cannot reach. Both are normal shapes for a question about access, where the point is to compare against what exists rather than to redesign it. Making those questions expressible is the methodological contribution of this work.

3.5 Spatial electrification layer

The layer classifies each community of the region as connectable, meaning that reaching it with a medium-voltage line costs less per household than equipping it with a standalone system, or dispersed, meaning the opposite. The classification rests on a cost comparison rather than on a fixed distance threshold, and no community is evaluated in isolation. The network is built as a single tree growing from the medium-voltage lines that already exist, so a line built to reach one community becomes available to the communities behind it. The cost of reaching a given community therefore depends on which other communities are already connected, and the layer builds the network and classifies it at the same time, following the greedy logic that spatial electrification tools apply at national scale [24].

At each step the cheapest remaining community is attached to the network as it currently stands: on an existing line, on a connected community, or by a mid-span tap on a line already built. A tap creates a junction node, which carries no household, no transformer and no connection. The candidate is accepted if its cost per household beats the annualised cost of a standalone system, and rejected

definitively otherwise, so the procedure terminates. For a candidate at a marginal distance d (km) from the network as it stands, with N households, the overnight cost per household is

$$\text{ext}_{\text{HH}} = \frac{c_{\text{line}} \cdot d + c_{\text{tr}}}{N} + c_{\text{conn}} \quad (15)$$

where c_{line} is the cost per kilometre of medium-voltage line, c_{tr} the cost of the transformer station that every connected community requires, and c_{conn} the connection cost per household. The result is annualised over the lifetime of the line. The breakeven is therefore endogenous. It varies with community size, since a larger community spreads the same line over more households. It also varies with the order in which the communities are treated, since d is measured against the network as it stands and not against the network that exists before the procedure starts.

The first community of a chain justifies the whole segment that opens the chain, and the communities behind it only pay their own branch. This is a sorting criterion and not a billing rule. It is also how a real network extends, since a line is justified by the first settlement it reaches and makes the next ones cheaper.

A greedy procedure carries no guarantee of optimality, since the order of examination shapes the tree, and the trees published here are 7.6 to 8.3 percent longer than the minimum spanning tree over the same nodes. Distances are straight lines, with no detour for rivers, terrain or rights of way, and a tap sits at its geometrically ideal point on a segment, which no built line would match. The cost is incomplete too. It carries medium voltage only, without the low voltage network inside each community, without any reinforcement of a conductor feeding a long chain, and without operation and maintenance. Edges never cross a cluster boundary either, so a community is always attached inside its own cluster even where a neighbouring line passes closer. The omissions push the cost down and the greedy tree and the cluster rule push it up. The two are not weighed against each other here, so the sign of the net error is not established.

3.6 Projecting the system to 2035 and 2050

The chain described so far produces a snapshot. Extending it to a future horizon requires four further quantities: the population and the number of households, the share of them that has electricity and by what means, the demand that follows, and a convention for reporting the cost of a system that inherits what an earlier horizon built. This section states the rules used to obtain them. Chapter 4 applies those rules to the study area and gives the resulting values.

3.6.1 Projecting the population and the number of households

The official municipal projections of the Bolivian statistical institute stop in 2034 [31]. Nothing official exists beyond that date at municipal resolution, and nothing official exists at all for the number of households, at any horizon. Both have to be built. The energy system is sized by households rather than by people, since a household is what receives a connection and what owns a refrigerator, so the link between the two carries most of the method.

The projection combines two rules. The regional total is projected by a deceleration rule anchored on the national growth profile, which the institute publishes to 2050 [32]. The growth rate of the region in year t is its 2034 value multiplied by a factor $\rho(t)$, which is the national rate of that year divided by the average national rate over 2032, 2033 and 2034. The three year average is needed because the national rates are of the order of 0.13 % per year by then, and a ratio of numbers that small is unstable without smoothing. The distribution of the regional total across the municipalities is projected by a

proportional logistic model calibrated on the relative trajectory of the official municipal shares. This is the architecture of the institute itself, a controlled total distributed by a share model [33].

One precaution governs every number obtained this way. The official projections start from a conciliated population, corrected for census undercount, which is higher than the census count itself. In 2024 the institute gives 11,916,453 inhabitants for Bolivia against 11,365,333 counted, a difference of 4.85 %. Using those levels directly would make the projected years inconsistent with the census baseline of the 2025 scenarios. Only ratios of the official series are used, applied to the census counts, so the levels of this thesis stay census levels throughout.

Household size is not projected directly. What is projected is the gap between the growth rate of households and that of the population, because that gap is the mechanism at work and it is directly observable on two intercensal periods. That gap narrows over time, and it is extended past the last census with a decay factor estimated on those two periods. Two starting gaps are used, one for the urban municipalities and one for the others, while the decay factor is shared. Section 4.10.1 gives the values and says which alternatives were rejected.

3.6.2 Projecting who has electricity, and by what means

The households are then split between the three populations of Section 4.2.1. The access rate of each municipality is projected first, as the share of dwellings with electricity from any source. That access is split between a line and a standalone system next. Both are finally applied to the household counts.

The order is forced by the data. The 2001 census records only whether a dwelling has electricity, not where that electricity comes from, and the categories for a private generator, a solar panel and other sources appear for the first time in 2012 [1]. The only variable available on three census dates is therefore the aggregate access rate, Source A and Source B together. The split between the two cannot be projected from the census at all and has to come from elsewhere. This is a data constraint rather than a methodological preference, and it explains the whole architecture of the method.

Access is projected by closing a constant fraction of the remaining gap to full coverage each year. Writing $a_m(t)$ for the access rate of municipality m in year t , that fraction is read from the last two censuses:

$$r_m = 1 - \left(\frac{1 - a_m(2024)}{1 - a_m(2012)} \right)^{1/12} \quad (16)$$

A constant closure approaches full coverage without ever reaching it. Since a trajectory is defined by the year in which it arrives, the closure is normalised over the length of the scenario:

$$\varphi_m(t) = \frac{1 - (1 - r_m)^{t-2024}}{1 - (1 - r_m)^T} \quad a_m(t) = a_m(2024) + \varphi_m(t) (1 - a_m(2024)) \quad (17)$$

where T is the number of years to the target. This reaches full access exactly at the target year without forcing the last value, and it preserves the relative positions of the municipalities. Where the estimated fraction is zero the closure is linear, $\varphi_m(t) = (t - 2024)/T$.

The rate r_m is estimated on the 2012 to 2024 period alone rather than on the full 23 years, and Section 4.10.1 says why, and what a straight line in percentage points per year gives instead.

No parameter is set by hand for any municipality. The rate r_m is estimated separately for each one and the form does the rest, so the pace is uniform in fraction of gap closed and heterogeneous in points

per year, which is the behaviour a real electrification programme has.

The second layer does not come from the census. It comes from the spatial layer of Section 3.5, which classifies every community as connectable or dispersed. For each municipality, the dispersed share is the fraction of the unconnected basin, Sources B and C together, that lives in a community classified as dispersed. That share is recomputed at each horizon, since the classification runs again on the projected household counts rather than being extended from 2024. A larger community spreads the same line over more households, so the extension cost per household of Equation 15 falls and more communities pass the breakeven: 626 in 2024, 645 in 2035 and 651 in 2050.

The third layer moves the households between the three sources, in three flows. Households already in Source A stay in Source A and grow with their municipality. Households in Source B split in two, since those in a dispersed community stay in Source B, no line ever being cheaper for them, while those in a connectable community move to Source A at the pace at which the line reaches them, measured by the fraction of the access gap already closed in their municipality. Households leaving Source C are shared between Source A and Source B according to the dispersed share of their municipality. The second flow is not a behavioural assumption, since the line is built to reach the Source C communities whatever happens and a Source B household on its path only pays a connection cost.

3.6.3 Projecting the demand

The sufficiency bundle is normative and fixed, so projecting it would only be a matter of counting households. No projection scenario carries it, as Section 4.10.5 states. The demand that has to be rebuilt is the measured demand of the connected households, estimated in Section 4.7.1 from utility sales. That method is reused unchanged, and three things about it change once the year is no longer 2025: appliance ownership evolves, the number of Source A households comes from the access trajectory above, and the calibration against the measured utility total disappears, because no measurement exists after 2025.

A data constraint decides how ownership is projected. The appliances that carry the demand have a single observation in time, since the census reports water heater, refrigerator, air conditioner, microwave and washing machine ownership for 2024 only. The appliances with a full history, radio, television, telephone, computer and internet, carry a few tens of kilowatt hours each. Everything with a time series is negligible and everything that matters has one point, so extrapolating a trend is not an option for the bulk of the demand.

What exists instead is variation in space. The 21 municipalities span the whole range of electrification in 2024, and that spread describes the path a municipality follows as it electrifies. Ownership is therefore regressed on the access rate across the 21 municipalities and projected forward with the access trajectory. Space is substituted for time, and the access projection feeds the demand directly. The substitution stays inside the observed range, since three municipalities already exceed 90 % access, so projecting a municipality towards full access interpolates between observed states rather than extrapolating beyond them.

The substitution was tested out of sample on the five appliances that do have a history. The relationship was fitted on 2012, used to predict 2024, and compared with the census, against a competitor that simply freezes ownership at its 2012 level. Three appliances out of five favour the access relationship and two favour freezing. The average advantage of 2.28 points is almost entirely carried by the telephone, which gained 13.3 points and whose penetration owes more to the rollout of the mobile network than to electrification. The two appliances that most resemble a refrigerator, durable

and permanently plugged in, are the two where freezing wins. The test therefore does not validate one rule for everything, and it is applied appliance by appliance, as Table 44 in Appendix D states.

Two further rules complete the demand. The scaling factor that calibrates the bottom up estimate on measured sales is frozen at its 2025 value, since it absorbs the appliances the census does not record, the gap between municipal ownership and the ownership of connected households, and the intensity of use, none of which evaporates in 2035. Only the residential sector carries such a factor. Public lighting and the general commercial tariff follow the number of Source A households, consistent with the allocation rules of Section 4.2.2, while the industrial tariff and the residual other category are frozen at their 2025 level, since they are Brazil nut, rice and timber processing and their scale follows harvests.

3.6.4 Costs, inheritance between horizons, and what may be compared

Four technologies carry a declining investment cost, taken from the moderate innovation scenario of the NREL Annual Technology Baseline [34]. The ratios of Table 2 are applied to the 2025 catalogue value at each horizon, never in cascade. Maintenance costs, construction emissions and lifetimes are unchanged, and every other technology keeps its 2025 cost, including diesel, the medium voltage line and the solar water heater.

Table 2: Investment cost ratios applied to the 2025 catalogue value. Dimensionless, from the NREL Annual Technology Baseline, moderate scenario.

Technology	2025 value	2035	2050
Utility solar	882.21 €/kW	0.600	0.458
Lithium battery	425.88 €/kWh	0.757	0.555
Household panel	2,730 €/kW	0.65	0.46
Household battery	594 €/kWh	0.76	0.61

The baseline has a 2022 base year, so these are ratios between two years of its series rather than reductions measured from 2025. Three of the four match the reductions the baseline publishes for the corresponding technology to within half a percentage point, which confirms that each ratio was paired with the right technology. The fourth, the lithium battery, has no directly comparable published figure. The ratios are dimensionless, which would protect them from the mixed vintages of the catalogue only if the base year of each matched the vintage of the value it multiplies, and it does not. The solar collector carries no ratio at all, for want of a source, so its price stays flat while photovoltaic capital falls to 46 percent of its value by 2050. Section 5.9 shows the water heating lever falling from 10.6 percent of the annual cost in universal sufficiency to 0.3 percent in 2050. That collapse has two causes, a hot water demand eight times smaller in the projections and a collector price that stays flat while the panel falls, and Section 6.5 separates them. The second cause is an artefact of the catalogue and must not be read as a statement about solar thermal collectors.

Each horizon starts from what the previous one installed. The capacity of 2035 becomes the brown-field floor of 2050, and Equation 8 charges no capital on it. The fleet of 2025 does not reach that far, since the diesel plants and the existing solar plant come to the end of their life before 2050, so their floor drops to zero at that horizon and a scenario that still needs diesel in 2050 has to pay for it.

Chaining the horizons this way breaks the brownfield rule. In 2050 the floor contains everything 2035 built, which is still being paid off, and the optimisation treats it as free, so a trajectory that invested

heavily in 2035 looks cheaper in 2050 for no real reason. The costs reported for the projections are therefore reconstructed. Capacity inherited from 2035 is re-invoiced at its 2035 cost, while capacity older than 2025 stays free because it genuinely is amortised. The reconstruction changes nothing at 2025 or at 2035, where the two figures agree to the euro, and it changes the 2050 runs only, always upward. The correction is proportional to what a trajectory built in 2035, which is why it cannot be left out: on the raw figure, building early buys a discount that nothing in the real system pays for.

That convention fixes what may be compared. The reconstruction re-invoices the fleet chained from 2035 but not the one that predates 2025, so the two horizons do not sit on the same accounting basis. Scenarios may be compared with each other at a fixed horizon. The change of one trajectory between 2035 and 2050 cannot be read causally, since it mixes a larger demand, cheaper technologies, the retirement of the old fleet and the completion of access.

4 Case study: the Bolivian Northern Amazon

Chapter 3 describes the two models and the spatial layer. This chapter applies them to the BNA, and most of it is data work rather than method. It starts from the territory and from the 2024 census, which splits the households of the region into the three populations used throughout the thesis. It then defines the five scenarios that Chapter 5 costs, groups the 21 municipalities into five clusters, and adapts the RAMP inputs to the climate and the economy of the Amazon. The two demand definitions are built next, the normative sufficiency bundle and the demand measured today. The next two sections configure EnergyScope cluster by cluster and apply the spatial layer that classifies the 697 communities of the region as connectable or dispersed. The last one builds the inputs of the two future horizons.

4.1 Regional context and study area

The BNA is a tropical lowland region of 21 municipalities, spread over the department of Pando, the north of Beni and the municipality of Ixiamas. It holds 317,517 inhabitants and 84,209 households. Cobija, Riberalta and Guayaramerín concentrate most of that population, and the rest of it lives in communities of a few dozen households each. The regional economy rests on the Brazil nut, on rice, on timber and on cross-border trade with Brazil and Peru. The Brazil nut is the dominant activity of Pando, and it is transformed in industrial plants in Riberalta and Cobija rather than in the communities that harvest it.

Distance is the constant of the region. Roads are few and heavy rainfall makes many of them impassable between November and April [2], so part of the territory is reached by river or by air for several months of the year. Everything the energy system consumes arrives the same way, diesel included, and the delivered cost of that diesel is well above the subsidised price at the pump. Cooking is done with firewood and bottled gas, and it is the largest household energy use of the region.

Electricity follows the same geography. The systems that exist are isolated diesel microgrids, and the public service covers 95.9 % of the households of Cobija against 0.3 % of those of San Pedro. A quarter of the households of the region are supplied by their own solar panel or generator, or by nothing at all, and no measurement records what those households consume. This is the situation the rest of this chapter describes quantitatively.

The national interconnected grid (SIN) does not reach most of the study area. An extension towards Pando has been formally planned since 2012: the Plan Óptimo de Expansión del SIN (CNDC) and the Plan Eléctrico 2025 (VMEEA) include a 1,025 km, 230 kV transmission line from Trinidad to Cobija

via Guayaramerín, Riberalta and Puerto Rico [35]. However, more than twelve years later, the project remains in pre-investment phase: no financial agreement has been secured and no construction has started.

In July 2024, Bolivia and Brazil signed a Memorandum of Understanding on electricity cooperation covering Guayaramerín, Riberalta and Cobija [36]. It is a framework for cooperation, with no volume, no price, no date and no duration, and it is not a power purchase agreement. A government press release announced an estimated saving of 69 million USD per year and 133,000 tonnes of CO₂ avoided annually, but those two figures come from the release rather than from the text of the agreement.

The supply is usually described as coming from the Brazilian hydropower plants of Santo Antonio and Jirau. Nothing reserves it. The loan proposal that would finance the interconnection states that ENDE Brasil would buy on the Brazilian free contracting market, and that public and private companies trading there could supply the energy destined for Bolivian demand [37]. Buying on an open market is not the same as buying from a named plant, so which plants would actually serve the region is undecided. Section 5.9 returns to the point, because the emissions benefit of the import depends entirely on that origin.

The agreement signals political will to find an alternative to extending the national grid, but the lines it would need have not been built. It foresees a first phase of medium-voltage connection and a second phase of high-voltage interconnection, both to be executed by ENDE. Until these are operational, the three municipalities remain isolated diesel systems.

The study area was defined from the electrical structure of the territory. It gathers the northern municipalities where the main interconnected grid is absent, or only marginally present near the boundary, and where access to electricity stays limited over large parts of the territory. Figure 3 shows the northern network physically separated from the central grid, with several independent microgrids inside it. The area therefore includes the whole department of Pando together with Ixiamas, Riberalta, Guayaramerín, Reyes, Santa Rosa and Exaltación. These municipalities also belong to the northern Amazonian lowlands and share the same geographic conditions, the same low accessibility and the same logistical constraints on electrification.




Figure 3: Delimitation of the BNA study area.

4.2 Census 2024 baseline

The demographic baseline comes from the official *Censo de Población y Vivienda 2024* of Bolivia [1]. The published tables were consolidated into one dataset covering the 21 municipalities of the study area, for the census years 2001, 2012 and 2024. It carries the source of electricity of the dwellings, the cooking fuel, the ownership of the appliances that drive the demand, vehicle ownership and access to information technology. It also carries the number of persons by electricity source, which gives the count of non-electrified people directly rather than through an average household size.

The differences between municipalities inside the region are very large. Figure 4 gives the three populations of every municipality as a share of its households. Public grid coverage runs from 95.9 % in Cobija to 0.3 % in San Pedro, and the same spread appears on the equipment: 73 % of the households of Cobija own a refrigerator against none in San Pedro. The figure also shows that the households supplied by their own panel or generator are not distributed like those with no supply at all, which is why the two are kept apart throughout this thesis.




Figure 4: Households by source of electricity in each of the 21 municipalities, as a share of the households of that municipality, from the Census 2024. Municipalities are ordered by public grid coverage. The total number of households is given on the right of each bar.

This heterogeneity is one of the main reasons for working at the municipality level rather than treating the whole region as a single block, as was done for the Lowlands in the national model [3].

The Census 2024 [1] reports the electricity source for each occupied private dwelling (*vivienda particular ocupada*). Collective buildings (hotels, hospitals) and dwellings that are unoccupied, for sale,

under repair or abandoned are excluded, since they do not generate residential energy demand. Table 3 shows the breakdown for the 21 municipalities.

Table 3: Households by electricity source in the BNA study area (Census 2024).

Electricity source	Households
Public grid	63,495
Own generator	3,766
Solar panel	4,307
Other	1,252
No access	11,389
Total	84,209

Three quarters of the households of the region are supplied by the public service, and the remaining quarter splits between an own panel or generator and no supply at all, as Figure 35 in Appendix E shows.

The total population across the 21 municipalities is 317,517 persons. The sufficiency demand covers all 84,209 households. The current demand treats them differently depending on how they are supplied today.

4.2.1 The three populations of the region

The electricity source reported by the census splits the households into three groups.

Source A, Grid-connected consumers (63,495 residential + commercial, industrial and public lighting).

All consumers connected to the public electricity service, including both the national interconnected grid and isolated microgrids. Their demand is estimated from measured utility data.

Source B, Off-grid electrified consumers (9,325 HH).

Consumers with electricity from a private source: solar panels (4,307), own generators (3,766) or other (1,252). Their demand is modelled with RAMP using reduced appliance bundles that reflect the limited capacity of autonomous systems.

Source C, Non-electrified households (11,389 HH).

Households without any electricity access. Their demand is zero by definition in the current demand.

Source A is the only group for which measured data exists. Sources B and C together represent 20,714 households, close to a quarter of the region, and they are the population this thesis is about.

4.2.2 Infrastructure allocation rules

The number of infrastructures modelled in each municipality follows allocation rules taken from the Lowlands column of [3], applied to the population and the household count of the Census 2024.

Schools and health centres are allocated at one per 5,000 inhabitants, with a minimum of one per municipality. Every other category follows the household count, also with a minimum of one. Table 37 in Appendix B gives the rules and the resulting totals.

The model only includes large schools and health centres, and this follows directly from the allocation rule. At one facility per 5,000 inhabitants, each one serves a large population, so it is configured with a full set of electrical equipment, including a cold chain for the health centres, rather than as a village school with a few light bulbs. Modelling smaller and more dispersed facilities requires sub-municipal data on community locations and populations that does not exist for the region.

4.3 Scenario definitions

Five scenarios are built for the year 2025, and Section 4.10.5 builds four more at the horizons 2035 and 2050. The two families carry different names, so a scenario is never ambiguous, and the horizon is written wherever the two could be read against each other. The five scenarios below differ along two axes. The first is the demand: how much electricity each part of the population is assumed to need. The second is the supply: whether the model is allowed to build new generation or only to run what already exists. Table 4 summarises the combinations.

Table 4: The five 2025 scenarios. Source A is the 63,495 grid-connected households, Sources B and C the 9,325 off-grid and 11,389 non-electrified ones.

Scenario	Demand, Source A	Demand, Sources B and C	Supply
Current system	measured	as measured today	existing plants only
Supply transition	measured	as measured today	free investment
Universal access	measured	sufficiency bundle	free investment, Source A supply frozen
Access transition	measured	sufficiency bundle	free investment
Universal sufficiency	sufficiency bundle	sufficiency bundle	free investment

The current system scenario reproduces the system as it works today. The demand of the grid-connected households comes from measured utility sales, the demand of the off-grid households is modelled with reduced appliance bundles, and the households without any supply have no electrical demand at all. On the supply side nothing can be built: the optimiser dispatches the diesel plants and the few solar installations that exist, and pays for their fuel and their maintenance. The result is a costed picture of the present situation, and it is the reference against which the other four are read.

In the supply transition scenario the demand is exactly the same as in the current system, and the model is allowed to invest in utility solar and in battery storage. Nothing else changes: the existing plants stay, the exchange topology stays, and no household changes its consumption. This isolates one question. How much of the current cost is caused by the way electricity is produced rather than by how much of it is consumed.

The universal access scenario asks the question a planner actually faces. What does it cost, today, to bring a decent level of service to the 20,714 households that are badly served or not served at all, without changing anything for those who already have electricity. The demand of the grid-connected households stays at its measured level, and the off-grid and non-electrified households move to the sufficiency bundle described below. Cooking is unchanged, since it does not depend on being connected.

Holding the existing system in place is what makes the comparison meaningful. If the model were free to re-optimize the supply of the grid-connected households at the same time, the difference in cost against the current system would mix the cost of the new access with the savings from rebuilding what is already there, and neither could be read separately. The production of the existing plants is therefore fixed to what they deliver in the current system, using the mechanism described in Section 3.4. The difference in cost then measures the access alone.

The access transition scenario has the demand of universal access and the supply rules of the supply transition. The 20,714 households of Sources B and C receive the sufficiency bundle, the grid-connected households keep their measured demand, and the model is free to invest anywhere. Nothing is frozen.

This scenario is an analytical intermediate rather than a proposal. Universal access and the supply transition each answer one question by holding the other axis still, so neither of them can say what access costs once the supply has already been rebuilt. That price turns out to be different, and the difference is a result in its own right. This scenario is also what allows the distance between the current system and universal sufficiency to be split into three steps that add up exactly, which is done in Section 5.8.

In universal sufficiency, all 84,209 households of the region receive the same normative bundle of energy services, whatever their current situation [14]. This includes the households of Cobija that are already connected and may consume more than the bundle today. The scenario is not a forecast of what people will consume. It describes what the energy system would have to supply if every household in the region were brought to a minimum acceptable level of service, and it is the most demanding of the five in terms of the population it covers.

Three of these scenarios change the demand of households that currently have none, so their cost is not complete until the lines needed to reach those households are paid for. For every scenario the cost is therefore reported in two scopes: the generation cost, which is what the optimisation returns, and the total cost including grid extension, which adds the network built by the spatial layer of Section 4.9. The two are never mixed in the same figure.

The five scenarios above are snapshots of 2025. They measure the gap between the present situation and universal access, but they say nothing about how the region gets there. Section 4.10 carries the region to 2035 and 2050 to look at the trajectory itself.

4.4 Clustering into EnergyScope cells

RAMP produces a demand profile for each of the 21 municipalities. However, running EnergyScope MultiCell with 21 separate cells is expensive and not necessarily meaningful, since many municipalities share the same characteristics. The municipalities are therefore grouped into a smaller number of clusters, and each cluster becomes one cell with its own demand profile, renewable resource availability and technology options. The method follows the approach developed for the national Bolivian energy system model [7], adapted to the BNA.

4.4.1 Clustering variables

The clustering uses a set of socio-economic, geographic and energy-related variables. Table 35 in Appendix A lists them with their sources.

Four municipalities have physical access to the national interconnected grid (SIN): Ixiamas, Reyes,

Santa Rosa Beni and Exaltacion. Four municipalities have gas distribution infrastructure: Cobija, Riberalta, Guayaramerín and San Lorenzo. All other municipalities operate on isolated microgrids or off-grid systems, and none has access to gas. These binary variables were verified against the geoportal grid map [38] and the ANH pipeline infrastructure map [39].

The electricity demand variable is the measured per-capita consumption from utility data [5]. It is zero for Bolpebra, Ingavi, Santos Mercado, Exaltacion and Santa Rosa Pando, for the reason given in Section 2.1: private panels and generators are not metered. This does not affect the RAMP sufficiency demand, which is computed independently.

4.4.2 Weight calibration

The clustering uses K-Medoids with a hybrid distance combining weighted feature similarity and geographic proximity [40]:

$$d_{ij} = (1 - \alpha) \cdot d_{\text{features}}(i, j) + \alpha \cdot d_{\text{geographic}}(i, j) \quad (18)$$

where d_{features} is a weighted Euclidean distance on standardised variables and $d_{\text{geographic}}$ is the Haversine distance between municipality centroids.

Rather than fixing the weights manually, they were determined through a collective assessment. A questionnaire was distributed to seven people familiar with energy planning or the Bolivian context. Each participant was asked to distribute 100 points across the 12 variables (11 features plus geographic distance) according to the importance they attributed to each one for grouping municipalities that should share energy infrastructure. The final weights were computed as the arithmetic mean of the seven responses, normalised so that the total sums to 1.0. Appendix A gives the seven individual responses and Figure 29 the resulting weights.

The two infrastructure access variables carry the highest combined weight, 0.306, which reflects a consensus that grid connectivity decides most of what two municipalities can share. Geographic distance comes next at 0.126, enough to matter without dominating the feature-based signal, followed by population density at 0.123 and GDP per capita at 0.094. The renewable resource and climate variables receive the lowest weights, consistent with the homogeneous tropical lowland setting of the study area.

4.4.3 Number of clusters

The number of clusters was evaluated with the Elbow method, the Silhouette score and the Gap statistic, each computed over 100 iterations per value of k for stability, and Appendix A shows the three curves. The three metrics agree on $k^* = 5$, which was adopted. The Silhouette score at that value is 0.151 and the Gap statistic is 0.536. A modest Silhouette score is expected for 21 points with several variables of low variance, and the resulting clusters are geographically coherent, which is the primary requirement for energy infrastructure planning. Running the 21 municipalities as 21 cells is possible and would resolve the exchanges and the lines where they physically exist, but each cell needs its own calibrated files and the exchange matrix grows from 10 pairs to 210, so the cost is calibration rather than computation. No other grouping was run. The three largest effects of Chapter 5 come from the diesel plants of C3 and C5 and from the demand definition, and neither depends on how the other municipalities are grouped.

4.4.4 Clustering results

Table 5 lists the 5 clusters and Figure 5 shows their geographic distribution.

Table 5: Clustering results: 5 clusters for the BNA.

Cl.	Municipalities	N	Description
C1	Exaltacion, Reyes, Santa Rosa Beni, Ixiamas	4	SIN-connected southern municipalities. Rural, low density.
C2	Bolpebra	1	GDP outlier (5,329 USD/person) from cross-border trade. Isolated, no grid, no gas.
C3	Guayaramerín, Riberalta, P. Gonzalo Moreno	3	Northeast urban corridor. Two cities with gas distribution. Highest population density after Cobija.
C4	Bella Flor, Filadelfia, Ingavi, Nueva Esperanza, Porvenir, Puerto Rico, San Lorenzo, San Pedro, Santa Rosa Pando, Santos Mercado, Sena, Villa Nueva	12	Rural Pando interior. No grid, no gas. Dispersed population. Represents the bulk of unelectrified demand.
C5	Cobija	1	Departmental capital. Highest density (120 inhab/km ²). Existing microgrid and gas distribution.




Figure 5: Geographic distribution of the 5 clusters across the BNA study area.

Two municipalities form singletons. Bolpebra (C2) has an unusually high GDP driven by cross-border trade at the triple frontier with Brazil and Peru, and Cobija (C5) is the only dense urban centre of the region, with a mature microgrid and gas distribution. The three other clusters separate the study area along its structural axes: grid access in C1, urban gas infrastructure in C3, and rural Pando with neither in C4. All non-singleton clusters are geographically contiguous, so every EnergyScope cell corresponds to an area where shared infrastructure is physically feasible.

4.5 RAMP input adaptations

The lowlands inputs from the RAMP Bolivia repository [41] were built for a macro-regional scale and do not reflect the conditions of the BNA. Three modifications were made, and they apply to both demand definitions.

4.5.1 Thermal comfort: fans calibrated by temperature data

The original lowlands inputs give every fan a fixed daily usage time of 540 minutes, identical across municipalities and seasons, and include no fan at all in winter. This does not describe the BNA. Winter here stays warm enough to justify fan use, and temperature varies widely between municipalities.

Hourly temperature data for 2019 was downloaded from Renewables Ninja [42] for each of the 21 municipalities. For every municipality m and season s , the daily average number of hours above 26°C is $\bar{h}(m, s)$. Jay et al. [43] recommend fan activation from approximately that temperature in hot and humid climates. The daily fan usage time, in minutes, is then

$$t_{\text{fan}}(m, s) = \max \left(60, \min \left(600, 30 \cdot \text{round} \left(\frac{60\bar{h}(m, s)}{30} \right) \right) \right) \quad (19)$$

which is the number of hot hours converted to minutes, rounded to the nearest half hour and bounded. The upper bound of 600 minutes is 10 hours, less than the 13 hours that the 8h–21h window spans, because even on the hottest days part of that window sits below 26°C . That bound binds in 45 of the 84 municipality and season pairs, so above 10 hot hours a day the calibration no longer separates one municipality from another. The lower bound never binds, since the smallest value produced is 210 minutes.

The usage window was also widened, from 9h–20h in the original inputs to 8h–21h. Average temperatures at 8h and 22h are nearly identical across the study area, 24.9 and 24.8°C , so extending the window symmetrically around the hottest hours is justified. Household members are also awake until at least 21h, according to the lighting windows already in the model.

The temperature series behind that calibration carries a spatial pattern more than a seasonal one. Winter, from June to August, is slightly warmer than summer in most municipalities. This is normal in a tropical climate, where winter is the dry season and carries less cloud cover. The more useful pattern is spatial. The variation between municipalities within one season reaches 6.5°C in winter, while the seasonal variation inside one municipality stays between 3 and 5°C . Calibrating per municipality therefore matters more than differentiating by season. Ixiamas ends up with a winter `func_time` of 210 minutes against 600 minutes for Santos Mercado, as Figure 6 shows.




Figure 6: Daily fan usage time (minutes) by municipality and season, calibrated from Renewables Ninja temperature data. The threshold for fan activation is 26°C.

The same calibration was applied to schools and health centres, which carried fixed fan usage times and no cooling in winter. Both now read `func_time` from the same lookup table as households, so the differentiation is consistent across sectors. School windows were adjusted to the two shifts of Bolivian schools in the Amazon, 7h–13h and 14h–18h [44], instead of the single 9h–16h window of the original inputs. Health centres keep the 8h–21h household window, since they operate throughout the day. They also carry two air conditioning units of 1 kW each, used one third of the fan lookup value. This reflects selective use in a rural health post, in the consultation room during visits and for medicine storage during peak heat, rather than continuous operation. The resulting 70 to 200 minutes per day is consistent with the fixed 120 minutes of the original inputs, with spatial and seasonal variation added.

4.5.2 Flour processing replaced by rice processing

The flour processing activity of the original lowlands inputs uses four machines, a grain dryer, a grain miller, a grain toaster and a pump, for approximately 26 kW of installed power. This activity belongs to the highlands and the valleys rather than to the productive reality of the BNA. It was replaced by a rice processing unit.

Rice is the main locally transformable cereal in the region. The dominant economic activity in Pando is castaña, the Brazil nut, but its transformation takes place in large industrial *beneficiadoras* in Riberalta and Cobija rather than at community level. Whether small-scale community processing could develop remains an open question.

The rice processing unit uses five machines, specified from Kapur et al. [45] and from manufacturer data, and Table 39 in Appendix B gives their ratings.

Total daily consumption is approximately 43.7 kWh, or about 22 kWh per tonne for a unit processing 2 tonnes per day. The literature reports 20 to 30 kWh per tonne for small-scale rice milling [45].

Rice is harvested from November to February and processed from December to March [46]. The model therefore runs rice processing at full capacity on weekdays in summer, at reduced activity in fall with operation roughly one day out of three, and not at all in winter or spring. This seasonal pattern lowers the contribution of income generating activities compared with the original flour processing, which ran all year at much higher power.

4.5.3 Radio power correction

The original inputs set radio power at 36 W, a value that probably dates from older models where radios carried integrated speaker systems. Off-grid radios sold in rural markets today are multi-function solar devices combining a radio, an LED lamp, a USB charger and sometimes Bluetooth. Product specifications from manufacturers active in off-grid markets [47, 48] put the audio function alone at 1 to 5 W, and 3 W is a reasonable upper estimate for a family radio with a built-in speaker. The lamp and charger functions are already modelled separately in RAMP, in the illumination and ICT inputs, so only the audio function needs to be counted here.

Radio power was therefore corrected to 3 W in the household, store and workshop inputs, and to 5 W for health centres, where a waiting room justifies a slightly more powerful unit. School inputs already used 4 W and were left unchanged.

4.6 Sufficiency demand

The sufficiency demand represents a normative energy access target rather than a projection of observed consumption. All 84,209 households in the study area are modelled with the same minimum-service appliance basket defined by RAMP, regardless of their current electrification status. This means that households already connected to the public grid are also included: some may currently consume above the sufficiency level (with inefficient appliances or higher usage), while others may consume below it. It does not attempt to reproduce actual consumption patterns; instead, it defines the demand that the energy system would need to supply if every household were to receive a minimum acceptable level of electricity services [14].

Table 38 in Appendix B summarises the household-level appliance inputs. These are inherited from the lowlands configuration of [3], with the adaptations described in Section 4.5 (thermal comfort, rice processing, radio power). All municipalities share the same appliance definitions, apart from the

fan usage time and the water heating of health centres, so a semi-urban municipality such as Cobija carries the usage patterns of a very rural one such as Ingavi. The refrigerator matters most here. It is described with three duty cycles in the warm seasons and two in winter, the same everywhere, although the temperature data needed to differentiate it is available, and those refrigerators draw 83 GWh in universal sufficiency, two thirds of everything delivered to an end use in that scenario. Its duty cycle is the assumption to which the size of the whole system is most sensitive, and a measurement campaign on a few dozen households is worth more to a follow-up study than any refinement of the optimisation.

Water heating is included in the sufficiency demand because it is part of the normative definition of minimum decent living conditions: access to warm water for basic hygiene. The water heating model calculates the power required to heat water from groundwater temperature to 40°C [8]. In the BNA, where the annual mean temperature is approximately 27°C, the resulting demand is moderate but still produces visible peaks in the morning (6h–9h) and evening (18h–20h) profiles (Figure 7).

However, Census 2024 reports that only 2.4 % of households in the region own a water heater, and 3.4 % in Cobija, even among those with electricity. This has to be read with care. Question 18.8 asks whether the household owns a *calefón* or a *termotanque* [49], a gas heater and a storage tank. No other item of that question covers water heating, so the instantaneous electric shower head, a device of about 4 kW that is widely sold across Bolivia and Brazil, leaves no trace anywhere in the questionnaire. The low figure therefore measures the ownership of two classes of equipment rather than the absence of water heating, and the share of households that heat water by some means is likely higher. What the figure does support is that few households own a dedicated water heating appliance, which is the basis for removing water heating from the current demand (Section 4.7).

For health centres, water heating is included in both demand definitions, since it serves a medical function (sterilisation, instrument cleaning). The groundwater temperature formula was corrected to use the annual mean air temperature plus a small attenuated oscillation (amplitude 1°C, phase lag 30 days), consistent with the thermal inertia of shallow tropical aquifers [50, 51].

Figure 7 shows the demand of Santos Mercado over a summer day, broken down by service. Cold storage runs as a constant baseload through the day and the night. Water heating produces two sharp peaks, in the morning between 6h and 9h and in the evening between 18h and 20h, and it is what sets the peak of the whole profile. Fans add a broad layer during daylight hours, rice processing appears during working hours, and ICT peaks in the evening. These peaks exist in the RAMP output only. Section 3.3.4 explains that the heat layer of EnergyScope carries the hot water demand as a flat profile, so they do not size the supply in any scenario. Section 5.7 quantifies what the choice of water heating technology is worth in energy and in cost.




Figure 7: Composition of electricity demand over 24 hours for Santos Mercado, sufficiency demand, summer day, all households. Water heating dominates the morning and evening peaks.

4.7 Current demand

The current demand aims to represent current electricity demand as closely as possible. The sufficiency demand applies one normative service target to every household. This one does the opposite: it estimates each of the three populations of Section 4.2.1 separately, because the data available for each is different. Source A is estimated from measured utility sales. Source B is modelled with RAMP, using reduced appliance bundles that reflect the limited capacity of an autonomous system. Source C has no electrical demand by definition. The three subsections below describe each method.

4.7.1 Source A: grid-connected demand from utility data

The demand of the connected consumers comes from the Anuario Estadístico de Electricidad 2024 of the AETN [52]. It reports electricity sales in MWh by tariff category, residential, general, industrial, mining, public lighting and other, for each isolated system. It gives no breakdown by end use. The measured total is therefore used as a calibration constraint, and the split by end use is built separately.

The regional total, about 180.1 GWh, is spread over the 21 municipalities in three ways, depending on what the yearbook publishes. Eight municipalities are read directly. Six others are served by the ENDE Cobija system, which the yearbook reports as one block of 70,571 MWh, and they are separated with a per household consumption borrowed from a comparable municipality, then scaled back so that the six sum to the measured block. Six municipalities are absent from the yearbook and receive the same borrowed rate, applied to their connected households from the census. The perimeter of two municipalities was adjusted. Guayaramerín includes its sub-systems of Cachuela Esperanza and Rosario del Yata, and excludes Puerto Ustárez, which lies outside the municipality. Three municipalities with an incomplete series are extrapolated on their own trend.

The residential total of each municipality is then split by end use, from the bottom up. The ownership rate of each appliance comes from the Census 2024, counted as Tiene over Tiene plus No tiene, with non-responses left out of the denominator. That rate, multiplied by the number of connected households and by an annual unit consumption, gives a first estimate, which is then scaled uniformly

to the AETN residential total. The measured total therefore decides the level, and the census decides how that level differs from one municipality to the next.

The unit consumptions come from two sets of sources. The refrigerator at 458 kWh/yr, the air conditioner at 451 kWh/yr and the electric water heater at 824 kWh/yr, with an electric fraction of 97.5 %, come from the MHE-GIZ characterisation study [53] and from Fundación Bariloche [54]. The television, the radio, the phone chargers, the laptop, the blender, the washing machine, the antenna and the lighting come from the RAMP parameters calibrated for the Bolivian Amazon, with residential lighting cross-checked against the UNEP LED project [55]. The fan combines its RAMP power rating with the thermal model of Section 4.5.1. The microwave, at about 50 kWh/yr, follows a Brazilian ownership survey [56].

Each appliance is then mapped onto an EnergyScope end use. The refrigerator goes to FOOD_PRESERVATION, the air conditioner and the fan to SPACE_COOLING, the electric water heater to HEAT_LOW_T_HW, interior lighting to LIGHTING_R_C, and everything else to ELECTRICITY.

The two other tariffs are treated with less detail, because no Bolivian end-use balance exists. The general tariff, which covers commerce and services, is split with the proportions of the 2016 Peruvian useful energy balance [57]. The industrial tariff goes entirely to mechanical energy, since the industry of the region is the small-scale processing of Brazil nut, rice and timber and no heavy industry operates there. Public lighting goes entirely to LIGHTING_P.

All these values are final energy, as the AETN measures it. The conversion to useful energy happens later, when the demand file is written, through the conversion coefficients of EnergyScope.

4.7.2 Source B: off-grid households modelled with RAMP

The 9,325 households that own a panel, a generator or another private source are modelled with RAMP, on the same framework as the sufficiency bundle but with a shorter appliance list. The sub-mix is 4,307 solar, 3,766 generator and 1,252 other.

The census gives appliance ownership per municipality, but as an average over all households, whatever their electricity source. Nothing in it says whether a refrigerator belongs to a connected household or to an off-grid one. The municipal rate is therefore applied to Source B as a first approximation, through the RAMP `occasional_use` parameter.

That approximation credits off-grid households with appliances their supply cannot run. A retention factor corrects it, appliance by appliance, from the power the supply technology can actually deliver. The factors are anchored on the Multi-Tier Framework [6], which grades access by available power and lists what each tier supports. A solar home system is Tier 2, above 50 W. It runs lighting, phone charging, television and a fan, but not a refrigerator, which draws power continuously and starts on a peak the system cannot follow. A generator is Tier 3, above 200 W, and runs the same appliances plus partial refrigeration. The other category is split evenly between the two tiers, which gives a mix of 53 % Tier 2 and 47 % Tier 3. The retention factor of an appliance is its weighted average over that mix, and the value written into RAMP is

$$\text{occasional_use} = \text{census ownership rate} \times \text{retention factor}$$

Table 6 gives the factors.

Table 6: Appliance assumptions for off-grid households (Source B).

Appliance	RAMP service	Ownership	Note
Interior lighting (4×7 W)	illumination	1.0	Universal
Exterior lighting (2×14 W)	illumination	1.0	Universal
Radio (1×3 W)	ICT	census	Low power, no reduction
Phone charger (4×5 W)	ICT	census	No reduction
TV (1×30 W)	ICT	census × 1.00	MTF Tier 2 appliance
Laptop (1×70 W)	ICT	census × 0.65	Tier 2/3 boundary
Refrigerator (200 W)	cold storage	census × 0.28	0 in solar, 0.6 in generator
Fan (1×30 W)	thermal comfort	census TV × 1.00	No fan data; TV as proxy

The refrigerator carries the whole effect. A generator household can run one, but not around the clock, so its Tier 3 factor is 0.6 and its weighted factor is 0.28. About 13 % of Source B households therefore operate a refrigerator in the simulation, against a municipal ownership rate near 48 %. Television and fan are Tier 2 appliances and stay at the full census rate.

The census does not record fan ownership, so the television rate is used as a proxy, as elsewhere in this work. Fans are probably more common than televisions in this climate, since they draw little power and answer a direct need, so the proxy understates fan use.

Water heating, air conditioning, microwave and washing machine are excluded. Water heater ownership is 2.4 % across the region. The three others either draw too much instantaneous power for an autonomous system or are rare off grid.

The infrastructure of Source B follows three rules of its own. Schools come from the RUE database of the Ministry of Education [58], restricted to active regular units and prorated where one district covers several municipalities. The off-grid count of a municipality is the school total multiplied by its off-grid household share, and those schools are entered as small schools, with lighting and basic ICT only. Health centres and public lighting are set to zero, since a health centre needs a reliable supply and public lighting is communal infrastructure already counted in Source A. Income generating activities keep the per household ratios of the sufficiency bundle, applied to the Source B household count.

RAMP produces the electrical profile alone. The diesel that a *motor propio* household burns to produce that electricity is accounted for on the EnergyScope side.

4.7.3 Source C: non-electrified households

The 11,389 households without any electricity access have zero electrical demand in the current demand.

Their cooking energy needs are accounted for separately in EnergyScope using the Census 2024 cooking fuel data, which provides the number of households by fuel type (wood, bottled gas, piped gas, electricity, other) for each municipality [1]. Households reporting electric cooking are excluded from this count since they contradict the “no electricity” classification. The cooking demand construction is described in Section 4.8.3.

4.8 EnergyScope configuration

4.8.1 Cluster-independent parameters

Table 7 lists the main parameters that apply identically across all clusters, and Table 43 in Appendix C the ones that differ between them.

Table 7: Cluster-independent model parameters.

Parameter	Value	Source / justification
Discount rate (<i>i_rate</i>)	10%	Article 48 of the Bolivian Electricity Law (Ley N° 1604, 1994) [59]; consistent with IRENA's rate for non-OECD countries [60] and World Bank practice for Bolivian rural electrification [61]
Network losses	15.96%	AETN 2024 [52]: ratio of gross production to sales for isolated systems in the BNA
GWP limit	1×10^{15}	Unconstrained (snapshot scenario, no CO ₂ target)
PV power density	0.2367 MW/km ²	National Bolivian value [30]

Resource prices come unchanged from the national model [30]. They are border prices and not subsidised end-user prices, the convention of the national model [30, 22]. Section 4.8.2 says how the fuels reach each cluster. The subsidised diesel price at the Bolivian pump, about 0.054 M€/GWh in 2025, is roughly half the border price used here, but hauling diesel to remote Amazonian communities partly closes that gap.

Three prices were recalibrated. The SIN electricity price is 0.0593 M€/GWh, from the CNDC wholesale price [62], since the national default of 5.064 M€/GWh was an uncalibrated European placeholder. The diesel border price of 0.1113 M€/GWh is kept. The LPG price is 0.0253 M€/GWh, from the YPFB subsidised garrafa price of 22.50 Bs per 10 kg converted at 6.96 Bs/USD. The conversion coefficients and the storage characteristics are unchanged, since every technology the BNA needs is already in the national database.

4.8.2 Inter-cluster exchanges

Distances between cluster centroids come from the Haversine formula, consistent with the national model [30, 22], and Appendix C gives the matrix. The shortest distance is 15.4 km, between Bolpebra and Cobija at the western edge of Pando, and the longest involve C1, which sits over 300 km from every other cluster.

Only cluster pairs that share a municipal boundary may exchange electricity, which leaves six pairs open at $t_{c_max} = 0.040$ GW and four closed. Three existing microgrid connections carry a forced minimum capacity, derived from Source A consumption with a load factor of 0.25: Cobija to Bolpebra at 0.111 MW, Riberalta to Villa Nueva at 0.247 MW, and Cobija to its five Pando satellites at 7.211 MW. The HVAC line costs 2 M€ per GW and per km, with maintenance at 4 % [22, 3]. Existing lines carry no investment cost, so only a new connection pays it.

All pipeline capacities are zero. Fuels enter as external resources delivered at their market price in each cluster rather than as inter-cluster freight, since YPFB distributes LPG and diesel through its own logistics network and the delivered price already includes transport.

4.8.3 Demand file

The demand file gives the annual end-use demand of each cluster, by sector and end-use category. It is built in two parts, the electricity-related demand from RAMP and the cooking demand from census data.

The minute-resolution profiles that RAMP produces are integrated over the year, converted to GWh, and mapped to the sector and end-use categories of EnergyScope [29], as Table 40 gives. The municipalities of a cluster are summed before the file is written.

Schools and health centres fall under SERVICES rather than HOUSEHOLDS, following standard economic sector definitions. Rice processing falls under INDUSTRY, as post-harvest transformation, and workshop machinery, a welding station of 5.5 kW and a grinder of 750 W, counts as commercial mechanical energy.

Cooking is the main household energy use in the BNA, and RAMP does not model it, since it is predominantly non-electrical. It enters the demand file as a fuel-independent useful energy demand of 1,344 kWh per household and per year, taken from the national model [30], and it applies to every household that cooks whatever the fuel. The Census 2024 [1] gives 82,501 cooking households across the five clusters, all fuel categories except *No cocina*. At 1,344 kWh each this is 110.88 GWh, and the restaurant kitchens modelled by RAMP bring the COOKING total to 111.00 GWh.

A useful energy demand rather than a fuel-specific final energy value avoids an important pitfall. Field measurements in the Oriente and the Yungas report 66.7 to 68.6 MJ per household and per day of final energy [63], but these zonal averages are dominated by firewood, whose stove efficiency is only 16 %. Applying such an intensity to LPG households overestimates their cooking demand by a factor of roughly four, since LPG stoves reach 64 %. EnergyScope takes the useful demand as input and converts it to final energy through stove-specific efficiencies, which handles the difference correctly. The dispatch that follows does not reproduce the fuel mix the census reports. A minimum capacity forces the wood stoves to be built but not to run, and nearly all cooking is served with LPG, which costs about three times less per unit of useful heat once the 16 percent efficiency of a wood stove is counted. The wood that a household gathers costs it nothing, and a cost optimisation cannot see that. The LPG volumes of Chapter 5 therefore describe a cost optimum and not the observed practice. The effect on the comparisons is small, since every scenario treats cooking the same way and cooking sits outside the electric scope.

The minimum installed stove capacity per cluster, for STOVE_WOOD and STOVE_LPG, follows from the census fuel counts and the useful demand:

$$f_{min} = \frac{\text{GWh}_{\text{useful}}}{(c_p \times 8760)} \text{ with } c_p = 0.1875$$

Space heating is zero, since a tropical Amazonian climate needs none. Passenger and freight mobility are zero as well, for lack of data at the municipal scale, and all mobility and freight technologies are excluded. Table 8 gives the resulting demand per cluster, which is what the optimisation reads.

Cooking and hot water together are 74 percent of the useful energy of the sufficiency bundle, and

Table 8: Annual useful energy demand by cluster and end use, for the sufficiency demand, in GWh/yr.

End use	C1	C2	C3	C4	C5	Total
COOKING	14.41	1.07	53.22	22.02	20.28	111.00
HEAT_LOW_T_HW	9.30	0.68	34.30	14.14	13.25	71.67
FOOD_PRESERVATION	3.84	0.28	14.16	5.83	5.47	29.59
SPACE_COOLING	2.23	0.19	10.05	3.99	3.52	19.98
ELECTRICITY	0.80	0.06	2.95	1.22	1.13	6.16
MECH. ENERGY_COMM	0.44	0.03	1.65	0.67	0.64	3.44
LIGHTING_R_C	0.22	0.02	0.82	0.33	0.32	1.70
MECH. ENERGY_IND	0.18	0.01	0.67	0.26	0.26	1.38
LIGHTING_P	0.03	0.00	0.10	0.04	0.04	0.22
Total	31.45	2.34	117.93	48.50	44.91	245.13

neither has to be electric. In the demand measured today, hot water is almost absent and cooking keeps its 111 GWh. Cooking alone is 111 GWh, more than all the other end uses put together. The choice of stove and of water heating technology therefore weighs as much as the choice of power plant, which is why every cost of Chapter 5 is split between an electric scope and a non-electric one.

Of the 62 GWh that remain, food preservation and space cooling take four fifths. A household of this region mostly needs to keep things cold, in a climate where nothing needs heating, and refrigeration runs through the night, so it is what sizes the storage rather than any evening peak in lighting or television.

Figure 8 puts the two demand definitions side by side, and they differ in shape and not only in level. Water heating is 29 % of the sufficiency bundle and almost nothing in the demand measured today, since few households own a water heater. Space cooling does the opposite, 8 % against 28, since the measured demand carries the air conditioning that the connected households own and the bundle does not. Cooking is nearly the same quantity in both, because every household cooks whatever its electricity supply is.




Figure 8: Annual useful energy demand by end use, for the two demand definitions. The area of each ring is proportional to its total. End uses below 3 % of both totals are gathered in Other.

C3 alone carries 48 % of the regional demand, and C3 with C4 and C5 carry 86 %. C2 is a rounding error at this scale, and the costs of Chapter 5 give the same picture.

4.8.4 Time series

The time series file contains 8,760 rows, one per hour, with normalised demand profiles and renewable capacity factors for each cluster.

For each cluster, the minute-resolution power of all its municipalities is summed, so that a municipality of 1,000 households weighs more than one of 200. It is then averaged to hourly resolution and normalised over the year:

$$\varphi(c, h) = \frac{\bar{P}(c, h)}{\sum_{h'=1}^{8760} \bar{P}(c, h')} \quad (20)$$

where $\bar{P}(c, h)$ is the mean power of cluster c over hour h , so that the 8,760 values of φ sum to 1. The result is a shape profile that tells EnergyScope how demand is distributed across the year, while the absolute level comes from the demand file. Two profiles are produced this way, ELECTRICITY for all RAMP categories combined and SPACE_COOLING for thermal comfort alone, since cooling has a daily and seasonal pattern the model needs to see separately.

Renewable capacity factors come from Renewables Ninja [42] for the year 2019, computed per municipality and averaged within each cluster. PV uses a 1 kW reference system tilted at the latitude of each municipality, between 11 and 14 degrees across the study area, and facing the equator, which is north in Bolivia. The mean capacity factor is about 17 %, typical for tropical regions. Onshore wind uses a Vestas V90 of 2 MW at 80 m hub height, which gives about 6 % and reflects the weak wind resource of the Amazon basin. Hydro, offshore wind, tidal and CSP carry placeholder values, since they are unavailable in the region.

Two gaps follow from this construction. The grid-connected consumers of Source A have no measured load curve, so the shape used for them is the one RAMP produces for the off-grid households of

Source B. Source B carries about 2 percent of the energy of the current system against 98 percent for Source A, and the seven projections reuse that shape unchanged, so any statement about peak demand or about storage sizing rests on that substitution. The capacity factors and the temperature series also come from the single year 2019, and no other year was tested, so the sizing of the solar and storage fleets carries the weather of that year with it.

4.8.5 Resources

Table 42 in Appendix C gives how resource availability was allocated to each cluster.

The region has no gas field, so local GAS availability is zero everywhere. Gas still appears in the results in small quantities, through the liquefied path: the model imports LNG, regasifies it and burns it in the existing gas water heaters of grid-connected households. Its catalogue price of 0.1988 M€/GWh is 1.8 times the diesel price and almost certainly reflects a European supply chain with liquefaction and shipping rather than a Bolivian one. It was not recalibrated, since the amounts stay below 0.6 M€/yr in every scenario and fall entirely outside the electricity scope.

4.8.6 Technologies

The technology file defines which technologies are available in each cell and their constraints. The base is the SA-PA 2025 configuration [30], with cluster-specific minimum capacities ($f_{\min}$) reflecting existing infrastructure.

The minimum diesel capacity of GENSET_DIESEL comes from the active power the AETN publishes for 2024 [52], 88.11 MW across C3, C4 and C5, and Table 9 gives it plant by plant. C1 and C2 have no local plant. Existing utility PV is present only in C5, 0.00510 GW at the Cobija solar plant.

The El Sena plant is a hybrid, and the yearbook reports its two halves separately: 5.10 MW of thermal capacity under ENDE, and 0.42 MW of photovoltaic capacity under ENDE Guaracachi, listed in a distinct section of the same table [52]. The diesel figure used above is therefore the thermal half alone and does not overstate the fossil capacity of C4. The photovoltaic half is a real installation, commissioned in 2018 and rated at 400 kW peak by its operator [64], but it is not entered as brownfield capacity in the model, so the existing solar capacity of C4 is understated by 0.42 MW.

The sufficiency demand includes domestic hot water (HEAT_LOW_T_HW, Section 4.6). In the decentralised heat layer, three technologies can serve this demand: an electric resistance heater (DEC_DIRECT_ELEC), a solar thermal water heater (DEC_SOLAR) and an electric heat pump (DEC_HP_ELEC). Each can be coupled to a hot water storage tank (TS_DEC). Their availability was configured as follows.

The solar water heater is present in the national technology catalogue but was not activated in the national runs. It is unlocked here, with the national cost parameters kept unchanged: 743 €/kW_{th} of investment, 8.25 €/kW_{th}/yr of maintenance and a 25 year lifetime. A dedicated solar thermal capacity factor series drives its hourly output. The roof area constraint is not binding: the installed collector area uses less than 3.5 % of the available surface.

This choice is well supported for the region. Heating water from about 27°C to 40°C is a small temperature lift, which is the ideal operating range for a solar thermal collector. A field trial in Manaus, in a climate comparable to the BNA, found that a simple thermosiphon delivered 150 litres above 45°C on at least 93 % of the days of the year [65]. A Bolivian supply chain also exists, since the El Sol project installed around 100 thermosiphon systems between 2009 and 2012 at about 1,414 USD for a 150 litre unit [66], the same order of magnitude as the model cost. The Brazilian Minha Casa

Minha Vida programme equipped social housing with rooftop collectors for around 896,000 people by 2015 [67]. A household with a thermosiphon on the roof is therefore a common configuration in the region and not a model artefact.

The storage tank uses the official EnergyScope TD parameters for decentralised thermal storage: 19 €₂₀₁₅/kWh of investment, 0.13 €/kWh/yr of maintenance and a 25 year lifetime [68], originally adapted from Rager [69]. Keeping these values preserves the consistency of the technology catalogue, since batteries, PV and gensets all come from the same database through the national model. The tank operates as daily storage with losses of 0.824%/h; energy cannot be carried over between typical days. Its construction emissions are not characterised in the database ($gwp_{constr} = 0$), so the hot water tanks carry no embodied carbon in any emission figure of this thesis.

The heat pump is kept in the catalogue but blocked ($f_{max} = 0$). On paper it is attractive, with a COP of 3, but it was not retained for the rural context of the study area. Heat pumps for domestic water are far less established than solar collectors or resistance heaters in the regional off-grid market, and their compressors and refrigerant circuits require a level of specialised maintenance and spare parts supply that is unlikely to be available in dispersed Amazonian communities. The realistic backup for a solar water heater here is the electric resistance element integrated in the tank, which is standard in the water heater products sold on the regional market. The optimiser is therefore left to choose between solar thermal collectors and direct electric resistance heating.

The minimum capacity of STOVE_WOOD and STOVE_LPG follows the rule of Section 4.8.3. Every other non-zero f_{min} is scaled from the SA-PA values in proportion to cluster population.

The distribution grid investment cost ($GRID_c_{inv}$) is set to zero. Existing distribution infrastructure is treated as a sunk cost. The cost of extending distribution lines to reach currently unserved households is not modelled by EnergyScope. It requires spatial data on household locations relative to the existing lines, which is the subject of the spatial layer (Section 4.9).

4.8.7 Current system configuration

The universal sufficiency scenario lets the optimiser build any technology in any cluster. The current system scenario represents the system as it exists today: a pure dispatch of installed capacity, with no new construction allowed. This requires a different technology configuration in EnergyScope, described below.

All generation and storage technologies are locked to their existing capacity, $f_{max} = f_{min}$, and to zero where nothing exists. The optimiser can dispatch existing plants but cannot invest in new ones. The brownfield correction of Equation 8 sets the capex on existing capacity to zero, so the cost is essentially fuel plus maintenance.

Table 9 gives the locked diesel capacities per cluster. For each plant, Cuadro VI-1 of the AETN yearbook reports both the installed capacity (*Potencia Instalada*) and the effective capacity (*Potencia Efectiva*); the *Potencia Efectiva* values are the ones used here. The ratio between the two is not constant across plants, which is why no single power factor reproduces the table.

This total is based on the AETN yearbook. The Geoportal spatial database used in Section 4.9.1 lists the same systems with plant-by-plant ratings that differ from the yearbook, since the two sources come from separate surveys taken at different dates; the Geoportal is used in this thesis only for the location of the plants, not for their capacity.

Existing utility PV is present only in C5, with 5.1 MW at the Cobija solar plant. Lithium-ion battery

Table 9: Existing diesel generation capacity by cluster (current system), from the AETN 2024 yearbook [52].

Cluster	Plants	Capacity (MW)
C1	none, supplied by the national grid	0
C2	none	0
C3	Riberalta 28.00, Guayaramerín 22.05, P. Gonzalo Moreno 3.28	53.33
C4	El Sena 5.10, Villa Nueva 2.32	7.42
C5	Cobija 27.36	27.36
Total		88.11

storage does not exist in the region, so it is set to zero. All other generation technologies are set to zero.

Demand-side and infrastructure technologies, such as stoves, refrigeration, lighting, HVAC lines and fuel storage, are left unconstrained so that the model can serve the demand. Lighting is restricted to LED, consistent with the fact that the demand was built using LED efficiencies. All transport technologies are set to zero, since mobility demand is zero.

The single electricity layer of a cluster does not distinguish a household connected to a line from one that is not, and Section 3.4 explains what was added to deal with it. The parameter `share_dispersion` is of no use here, since it can only push the production of the home systems up and what is needed is the opposite, so it stays at zero in all five clusters. `HS_DIESEL` carries instead a floor and a ceiling set to the same value in every cluster, so the annual output of the household generators is fixed rather than merely bounded, at the demand of the Source B households of the cluster and nothing more. `PV_HS` needs no bound, since its capacity is already fixed and its output follows the solar resource.

The existing off-grid equipment enters as brownfield capacity. `PV_HS` is fixed at the existing panels, rated at 50 Wp each [70], and `HS_DIESEL` at the existing generators, rated at 500 W each, a value assumed for lack of any field data for the region, and it sets the brownfield capacity credited to that technology. Its capacity factor c_p , the annual energy ceiling of Equation 14, is set to 0.85, where the catalogue carries three unsourced values for the same family of machines. Where the capacity is locked to the existing generators, that ceiling does not bind. Where the capacity is free, the ceiling sets the capacity the solver installs, since nothing more than the fixed production is needed. `BATT_HS` is zero, since the existing pico-batteries of 123 Wh per PEVD kit amount to roughly 530 kWh across the region, which is negligible at cluster scale. The constraint that forces the home battery to charge only from the home panel is deactivated here, since it would force `PV_HS` to zero where the battery capacity is zero.

The exchange topology is restricted to the three connections that physically exist today, C2 to C5, C3 to C4 and C4 to C5, and the three pairs open in universal sufficiency are closed. C1 remains the only cluster with access to the national grid, and its external electricity availability is uncapped here, so the model can serve the whole C1 demand from the SIN, which is what happens today.

4.8.8 Universal access configuration

The universal access scenario measures what it costs to bring the sufficiency bundle to the 20,714 households of Sources B and C while leaving the connected households exactly as they are. Section 4.3 says why the second half of that sentence is a modelling requirement rather than a convenience. Everything that serves Source A today is therefore held at its current level, through the annual production bounds of Section 3.4. This section gives the values and says which technology carries which bound.

The three clusters with a local diesel plant carry a floor on annual production, set to what those plants deliver in the current system. Table 10 gives the values. The three sum to 184.921 GWh, which is the entire regional genset output of the current system, so the counterfactual is reproduced exactly rather than approximately. The two clusters without a local plant carry no bound at all, since there is nothing to hold in place.

A floor alone is enough here, with no matching ceiling. The run confirms it: the existing fleet produces the same 184.921 GWh in universal access as in the current system, so the floor binds and nothing pushes these plants above it. The new demand is served entirely by what the model builds on top. This is not a general rule, since Section 5.5 shows the gensets running more once the same demand is added to a rebuilt supply.

Power plants are not the only equipment that Source A already owns. Some connected households heat water with a gas boiler and others with an electric resistance, and both appear in the current system. Left free, the optimiser would replace them with solar collectors, which is cheaper, and the saving would be counted as a benefit of giving access to somebody else. Two further floors prevent this. DEC_BOILER_GAS carries one in C2, C3, C4 and C5, and DEC_DIRECT_ELEC carries one in C1, C3 and C5, each set to the production of the same technology in the current system. Where a cluster shows a floor of zero in Table 10, the technology is genuinely unused there in the current system, and the zero is stored deliberately rather than left unset. The five gas floors sum to 2.433 GWh and the five electric floors to 2.669 GWh, which are the productions of the two technologies in the current system.

The home system technologies raise the same problem here as in the current system, so HS_DIESEL again carries a ceiling on its annual production.

The difference with the current system is that it carries no floor. In the current system the household generators were held to exactly today's output, because that scenario reproduces the present situation. Here the households of Sources B and C are being re-equipped, so there is no reason to hold their generators at any particular level. The ceiling states what the technology cannot exceed, and the optimiser chooses the mix of panel, generator and battery underneath it. The ceiling is set to the whole electricity demand of the dispersed households of the cluster, not to the share of it that a generator would plausibly cover. It is therefore a bound on what the household generators could deliver if they served that demand on their own, and C5 receives a ceiling of zero because no household there falls outside the reach of the grid. The same ceiling is written on the household panel, and it applies to each technology separately and never to their sum.

The spatial layer of Section 4.9 classifies each community as connectable or dispersed, and share_dispersion tells the model how much of the electricity must come from home systems. Here it is calibrated as an energy ratio, the electricity demand of the dispersed households of a cluster divided by its local electricity production, since the constraint is written on local production alone and excludes inter-cluster exchanges. Universal sufficiency calibrates the same parameter the same way, as Table 14 gives.

The spatial layer classifies all sixteen communities of C2 as connectable, at the three horizons. The cluster therefore has no dispersed household to equip, and both the floor and the ceiling on its home system technologies are set to zero. This is a result of the classification and not a modelling choice. C2 is also the cluster where the share constraint could not bind in any case, since it imports almost all of its electricity from C5 and the constraint is written on local production alone. Section 4.9.3 describes that weakness and Section 8.1 says how to remove it. It has no consequence here.

Table 10: Annual production bounds applied in the universal access scenario, in GWh/yr. The floors on the existing equipment reproduce its production in the current system. A dash means that the technology is not present in the cluster; a zero is a bound that is set and binding.

Technology	Bound	C1	C2	C3	C4	C5
GENSET_DIESEL	floor	—	—	103.859	25.653	55.409
DEC_BOILER_GAS	floor	0	0.0172	0.8824	0.6510	0.8824
DEC_DIRECT_ELEC	floor	0.3656	0	1.6132	0	0.6897
HS_DIESEL, PV_HS	ceiling	0.1794	0	0.1107	0.2419	0

Everything that does not serve Source A is open. The model may build utility solar and battery storage in any cluster, and it may equip the new households with home systems. The solar water heater DEC_SOLAR is unlocked, unlike in the current system where it is blocked because the connected households do not own one. This last point matters for the reading of the results, since part of the cost difference against the current system comes from a technology becoming available rather than from the access itself, and the two are separated in Section 5.7. The exchange topology is the same as in the current system, restricted to the three connections that physically exist, and C1 keeps its access to the national grid.

4.9 Spatial layer application

The EnergyScope optimisation in Section 3.3.1 treats the distribution grid as a sunk asset (GRID $c_{inv} = 0$) and assumes every household is already reachable. In practice, a significant share of households in the region has no grid connection. This section describes the spatial layer that estimates the cost of physically reaching those households, completing the generation optimum with the infrastructure dimension that the optimisation leaves out

The spatial layer feeds two quantities into EnergyScope: a `share_dispersion` per cluster, which tells the optimisation how much of the electricity of that cluster must come from standalone home systems, and a grid extension capex per cluster, which is the cost of building medium-voltage lines to the communities that are within economic reach. Together with the EnergyScope generation cost, these define the total system cost of universal access. The comparison between the generation cost alone and the total cost including grid extension is the core of Objective 3.

4.9.1 Data sources

Three GIS layers from the Geoportal of the Ministerio de Hidrocarburos y Energías are used.

- **Community locations** (`comunidades_electricidad_2012`): point layer with, for each community, the number of households without electricity (*No tiene*) and the number with private off-grid supply (*Motor propio, Panel solar, Otra fuente*). These correspond to Source C and

Source B of Section 4.2.1 respectively. The layer dates from 2012; the 2024 Census provides updated totals per municipality, but the spatial distribution within each municipality is assumed stationary. In practice the 2012 counts are used as relative weights: each community receives a share of the 2024 municipal total in proportion to the weight it had in 2012. The position of a community and its size relative to its neighbours therefore come from 2012, while the number of households it stands for comes from 2024.

- **Medium-voltage lines** (Líneas de Media Tensión 2023): LineString layer with 653 features, distinguishing interconnected (SIN) from isolated (Sistema Aislado) lines.
- **Generation plants** (Centrales Generadoras): point layer with the 22 plants of the region, giving about 89.6 MW of diesel and hybrid capacity and 5.1 MW of utility PV. This layer is used for the location of the plants only; their capacity comes from the yearbook, as explained in Section 4.8.7.

The distance from each community to the nearest medium-voltage line segment was computed using the line geometry, not the generation plant locations. Only cluster C1 has communities nearest to a SIN line; clusters C2 to C5 are entirely served by isolated mini-grid lines. Short extensions of the SIN are treated as feasible.

4.9.2 Cost-breakeven classification

The breakeven classification of Section 3.5 is applied here with the unit costs of the region.

The unit costs in Equation (15) are 8,316 €/km for the medium-voltage line, 3,234 € per connected community for the transformer station and 138.6 € per household for the connection, converted at 0.924 from the 9,000 USD/km, 3,500 USD and 150 USD of Peña Balderrama et al. [25] (Table A.5). The extension cost is annualised with $CRF(0.10, 30)=0.106$.

The standalone option is a PV plus battery home system, sized on the electricity E that a dispersed household draws under the sufficiency bundle, about 1,169 kWh/yr. That figure is the bundle without water heating plus the electric backup of the solar collector, which the optimisation puts at 4.83 percent of the hot water demand.

A thermosiphon solar water heater serves the water heating in both options, so its cost appears on both sides of the breakeven and cancels. Only the electric supply is compared. The panel is sized on the annual energy and the battery on the days of autonomy:

$$P_{PV} = \frac{E}{c_f \cdot 8760 \cdot \eta} \quad C_{bat} = \frac{n_{aut} \cdot E}{365 \cdot DoD} \quad (21)$$

with c_f the PV capacity factor of the cluster, η the end to end efficiency of the installation, n_{aut} the days of autonomy and DoD the depth of discharge of the battery. The two components carry different lifetimes, so each is annualised with its own capital recovery factor and the operation and maintenance rate applies to the whole:

$$sa_{HH} = CRF_{PV} c_{PV} P_{PV} + CRF_{bat} c_{bat} C_{bat} + \omega (c_{PV} P_{PV} + c_{bat} C_{bat}) \quad (22)$$

A community is dispersed when the extension cost of Equation (15), annualised on the same basis, exceeds sa_{HH} , and Figure 9 shows that comparison community by community. Table 11 gives the values. The result is 1,265.5 € per household and per year on average, and it runs from 1,212.80 to 1,284.99 across the 21 municipalities, since c_f differs between clusters.

Table 11: Parameters of the standalone system. The unit costs and the lifetimes come from the technology catalogue of Section 4.9.4, so the classification and the optimisation price the same equipment.

Parameter	Symbol	Value	Source
End to end efficiency	η	0.75	Panjon [71]
Days of autonomy	n_{aut}	2	Lighting Global [72]
Depth of discharge	DoD	0.80	Lighting Global [72]
Panel cost	c_{PV}	2,730 €/kW	Section 4.9.4
Battery cost	c_{bat}	594 €/kWh	Section 4.9.4
Panel lifetime		20 yr	Section 4.9.4
Battery lifetime		10 yr	Section 4.9.4
Operation and maintenance	ω	2 %/yr of capex	Peña Balderrama et al. [25]

The battery carries most of the standalone capex, so the days of autonomy is the assumption the classification is most sensitive to. Lighting Global gives 3 to 5 days for critical public facilities and 1.5 to 2 days once the load is not critical, which places a household at the two days used here [72]. Table 12 varies that value and three other levers, one at a time, and the published count of 626 sits inside the range they produce. Varying the depth of discharge between 0.70 and 0.90 is equivalent to 1.78 and 2.29 days of autonomy, which that range already covers. The last line of the table moves the line cost instead. A last mile electrification plan published for Paraguay in 2025 prices a 22.9 kV three phase line at 8,212 and 12,262 USD/km, poles included [73]. The 9,000 USD/km used here falls between the two, and taking the upper value, 11,330 €/km after conversion, leaves 578 communities connectable in 2024.

Table 12: Connectable communities under each assumption, one variant at a time. The variants are not cumulated. The last line raises the line cost instead of the standalone cost, and leaves the standalone system unchanged.

Variant	Standalone (€/HH/yr)	2024	2035	2050
Published	1,265.5	626	645	651
$n_{\text{aut}} = 1$ day	830.9	547	586	616
$n_{\text{aut}} = 3$ days	1,700.0	646	658	671
Battery lifetime 15 yr	1,116.8	603	636	646
$\eta = 0.85$	1,218.8	619	643	651
Line at 11,330 €/km	—	578	606	638




Figure 9: The breakeven, one point per community. The dashed line is the median annualised cost of a standalone system and the band its range across the 21 municipalities. The vertical axis gives what a community pays for its own spur to the existing network, so the connectable communities that sit above the band are the ones the shared trunk makes viable.

4.9.3 Source B and Source C integration

The analysis covers the 11,389 households of Source C and the 9,325 of Source B, whose locations are known from the 2012 community layer. Source B households are modelled at the full sufficiency level, since their panels and generators cannot deliver the bundle. A connectable Source B household adds a connection cost to the extension capex but no line cost, because the line is built for the Source C communities in any case. A dispersed one keeps its existing equipment, credited as brownfield capacity on PV_HS and HS_DIESEL at 50 W per panel and 500 W per generator. Those credits stay below 0.1 MW per cluster and change nothing in the extension result.

The parameter `share_dispersion` carries the classification into the optimisation. It appears in a constraint written for every hour of every typical day. On one side sits the output of the home system technologies, the panel and the generator plus what the home battery discharges net of what it absorbs. On the other sits the electricity supply of the cluster multiplied by the parameter. The constraint requires the first to be at least the second.

The definition of that supply matters most here. It is the output of every technology that feeds the electricity layer, plus the electricity that arrives as a resource, and it leaves out what a neighbouring cluster sends. Two consequences follow, and both matter later. The denominator is a supply and not a demand, so it exceeds what the end uses of the cluster receive, by the network and the storage

losses. A cluster that covers most of its consumption through an exchange also has an almost empty denominator, so the constraint cannot bind however high the parameter is set.

The parameter is not derived the same way in the two scenarios. Universal sufficiency uses a ratio of households, taken directly from the breakeven above. Universal access uses a ratio of energy, calibrated so that the constraint delivers the annual electricity demand of the dispersed households. Section 4.8.8 explains why the two cannot be the same. The share of universal sufficiency is larger in C1 and C3 and smaller in C4, as Table 14 shows, and the reason sits in the denominator. Under the bundle every household consumes the same amount, so the dispersed households weigh as much in the energy as they do in the population. That stops being true as soon as the connected households keep a demand several times larger than the bundle.

The constraint has two known weaknesses. Electricity arriving from a neighbouring cluster is not counted in the local supply, so a cluster that imports most of its consumption has an almost empty denominator and the constraint cannot bind; C2 is that case, and it happens to have no dispersed household, so nothing is lost here. The parameter is also calibrated on an annual total while the constraint is written for every hour, which leaves the home systems between 0.8 and 2.1 percent above their target in 2025 and between 101 and 129 percent of it in the projections, where the dispersed share and the annual ceiling interact non-monotonically. The ceiling is also written on each technology separately and never on their sum, so the household panel and the household generator together deliver about 2.4 GWh more than three clusters need in the projections, against a regional supply of 345 to 466 GWh.

The extension capex covers the lines built to reach the communities of Source C and the connection of the Source B households within economic reach. It does not cover the replacement of the existing equipment of the households that stay dispersed, although those households are modelled at the full sufficiency bundle, which their current panels and generators cannot deliver, so the total cost of universal access and of universal sufficiency is a lower bound.

Table 13 gives the classification and the cost of the network per cluster, and Table 14 the two sets of parameter values.

Table 13: Spatial classification and grid extension cost per cluster, at the household counts of the 2024 census, annualised with CRF(10%, 30 yr). These values are shared by every scenario that gives the sufficiency bundle to Sources B and C.

	C1	C2	C3	C4	C5	Total
Connectable communities	144	16	150	302	14	626
Dispersed communities	20	0	10	41	0	71
Tree length (km)	1,471.36	132.90	709.36	1,626.10	8.19	3,947.91
Extension capex (M€)	13.37	1.25	7.25	15.56	0.20	37.63
Annualised (M€/yr)	1.42	0.13	0.77	1.65	0.02	3.99

At the household counts of the 2024 census, 626 of the 697 communities are connectable and 71 are dispersed, for 3,947.91 km of medium voltage line over 755 edges. Among the connectable communities, 285 attach directly to the existing network, 212 to a community connected before them, and 129 by a mid-span tap, which creates a junction node carrying no household and no transformer. Connectable is a screening result and not a state of the network, since none of these lines exists today, and a dispersed community is served by its own panel and battery rather than left without electricity.

Table 14: The `share_dispersion` parameter passed to EnergyScope, by cluster and by scenario.

Scenario	C1	C2	C3	C4	C5
Universal access	0.0080	0.000	0.0010	0.0093	0.000
Universal sufficiency	0.0113	0.000	0.0014	0.0080	0.000

The line itself carries 87 % of the extension capital, against 5 % for the transformer stations and 8 % for the household connections. What decides the classification is therefore the distance to the network, and not the equipment at either end, as Figure 36 in Appendix E shows.

That distance is a marginal one, measured against the network as the procedure has built it so far. A single community can therefore justify a long segment that every community behind it then reuses, which is what makes a chain viable when no member of it would pay for the whole line on its own. Cluster C1 shows the effect at its clearest, in Figure 10.




Figure 10: Published tree of cluster C1. The three deepest paths share a trunk of 9 edges and 165.9 km, which one community alone justifies and the others reuse.

4.9.4 Standalone technology costs

The standalone technologies are PV_HS (panel, inverter, cabling), BATT_HS (battery, charge controller, cabling) and HS_DIESEL (small household generator). The cost of these technologies is uncertain: component prices vary between sources and change quickly. Two cost bases are given in Table 15.

The Arias values are built bottom-up from Bolivian component prices [3]. Inspection showed that the inverter (576 USD for a 1.2 kW unit) accounts for about 60 % of PV_HS and is not independently sourced in the original thesis. The 2024 values replace the inverter and charge controller prices with current retail listings from AutoSolar (Peru and Colombia) and cross-check the total against

Table 15: Standalone technology cost assumptions.

Technology	Arias [3]	2024 sourced	Source (2024)
PV_HS (€/kW)	4,747	2,730	Inverter: AutoSolar PE/CO; panel, cabling: retail
BATT_HS (€/kWh)	659	594	Lithium pack, charge controller and cabling, retail
HS_DIESEL (€/kW)	1,302	1,302	Arias [3] (unchanged)

the deployed household cost of the World Bank IDTR rural electrification programme [61], which reports about 677 USD per household for an SHS installation. The battery cost was updated with the BloombergNEF 2024 lithium pack benchmark. The model runs use the 2024 sourced values. The Arias values are given for comparison and no run uses them.

4.10 The 2035 and 2050 horizons

The sections above describe the region as it stands in 2025. This section applies the projection rules of Section 3.6 to it, and produces the inputs of the two future horizons: the population and the households, the split between the three sources, the demand of the connected households, and the scenarios that Chapter 5 runs.

Two trajectories are built. The first brings every municipality to universal access in 2035. The second reaches the same point in 2050, and its shape follows what the region has done since 2012 rather than what it has announced. Universal access means that Source C is empty, not that Source A covers every household, since a household supplied by its own solar home system has access to electricity and simply has no line. Some households therefore remain in Source B in 2050. They are the ones for which the spatial layer of Section 4.9 finds a standalone system cheaper than a line extension, and time does not change that arithmetic.

One property makes the two trajectories comparable. The population and the number of households are identical in the two, in both years, and only the split between Sources A, B and C differs. Any cost difference between them therefore comes from the pace of access alone, and never from a different assumption about how many people live in the region.

4.10.1 Calibrating and testing the projection rules

The rules of Section 3.6 carry parameters, and choosing between candidate rules requires testing them on the territory they project. This section does both for the BNA, and it is what selects the rules that the rest of this chapter applies.

The choice between candidate models follows the protocol the institute documents for its own municipal work: calibrate on the 2001 and 2012 censuses, predict 2024, and compare with what the 2024 census counted [33]. Six models were tested, the five the institute documents plus the deceleration rule of Section 3.6.1. The criterion is the error on the population of each cluster rather than of each municipality, because the clusters are what enters EnergyScope. The error on the regional total discriminates nothing, since it is zero by construction for any hierarchical model. Table 16 gives the result.

The two models constrained by a control total take the first two places and the three free models collapse, as Figure 33 in Appendix D shows.

Table 16: Backtest of the six candidate models. Each is calibrated on the 2001 and 2012 censuses and used to predict 2024. The figure is the deviation of the predicted cluster population from the 2024 census, in %.

Model	C1	C2	C3	C4	C5
Proportional logistic	−3.6	9.4	−7.2	2.9	5.9
Share of growth	−3.7	14.7	−8.2	11.5	13.7
Logistic	1.4	30.4	−3.6	28.5	31.0
Linear	1.2	31.0	−3.5	29.6	31.8
Deceleration rule	0.1	34.5	−5.3	50.3	41.9
Exponential	7.6	68.2	1.0	115.0	87.7

A free functional form fitted municipality by municipality gives a median absolute error near four times higher, 45.58 % against 11.92. Medians are used rather than means, because the mean is pulled by a few very small municipalities. The result should be read narrowly. A free form fails at municipal scale on a territory where several municipalities hold fewer than a thousand households and some reverse their trend from one census to the next, and it says nothing about the same form applied to large aggregates.

Table 16 places the deceleration rule fifth, which looks like a reason to drop it. Its failure mode was measured instead. Its municipal error correlates strongly with the 2001 to 2012 base growth rate, with a Pearson coefficient of 0.774 at $p < 0.0001$, against 0.220 and not significant for the proportional logistic. The rule propagates the error of the rate it is given, and its three worst municipalities are three of the four with the highest base rates. That failure mode does not apply where the rule is used, since the official municipal rates for 2033 and 2034 run from -0.155 to $+1.113$ % per year, an order of magnitude below the rates that break it. On the regional aggregate, where no control total applies, its backtest error is 6.32 percent against 9.48 for the free logistic, so it is kept for the total.

The gap between the growth rate of households and that of the population is calibrated on the same two periods. Between 2001 and 2012 the population of the study area grew by 2.920 % a year and the households by 4.607, a gap of 1.687 points. Between 2012 and 2024 the two rates are 1.420 and 2.503, a gap of 1.083 points. The gap therefore falls by a factor of 0.642 over twelve years, and it is extended past 2024 with that same factor. Two starting gaps are used, 0.934 points per year for the three urban municipalities and 1.373 for the other eighteen, while the decay factor is regional and shared. Estimating a separate factor per group was tried and rejected, since it rests on two points only and returned 0.981 for the rural group, which means no deceleration at all over 26 years.

The access rate is tested the same way. The alternative to the gap closure of Equation 16 was a straight line in percentage points per year. Both forms were calibrated on 2001 to 2012 and tested on their prediction of 2024. The gap closure gives a mean absolute error of 8.4 points against 18.5 for the linear form, and a bias of $+2.4$ against $+13.4$. Excluding two outlier municipalities the errors are 5.2 and 15.5 points. A second argument is more telling than either figure. The linear form sends 13 of the 21 municipalities to full access by 2024, and not one of them is there. A rate expressed in points per year cannot slow down as it approaches the ceiling, so it runs straight through it, while the gap closure form produces that slowdown on its own with no extra parameter and no manual cap.

Both forms overpredict 2024, and both do so in the same direction, because electrification was much

faster in the first intercensal period than in the second. Across the 21 municipalities the access rate rose from 48.6 to 78.9 % between the 2001 and 2012 censuses, then to 86.5 % in 2024, which is 2.76 then 0.63 points per year over the eleven and twelve years that separate them [1]. Leaving out the three urban centres gives 2.98 then 1.39, so the slowdown appears in both groups and is not an artefact of saturation in the cities. Calibrating over the full 23 years would project a pace that no longer exists, so only the 2012 to 2024 period is used.

4.10.2 Population and households in 2035 and 2050

The census baseline of Section 4.2 gives the starting point, 317,517 inhabitants and 84,209 households in 2024. The census was taken in 2024 and the scenarios of Chapter 5 are dated 2025, since the utility data they use covers that year. The two are treated as one starting state throughout, and the tables below carry the census year.

The correction for census undercount is larger here than nationally. The statistical institute gives 359,200 inhabitants for the region against 317,517 counted, a difference of 13.13 %, against 4.85 % for the country as a whole. Following the rule of Section 3.6.1, only ratios of the official series are used, so those 41,683 inhabitants never enter the levels of this thesis. The size of the correction still says something. The institute corrects the BNA almost three times harder than the country, and the households a census taker misses are likely to be the ones a medium voltage line misses too.

The deceleration factor $\rho(t)$ applied to the regional rate falls slowly and then collapses: 0.9719 in 2035, 0.9545 in 2040, 0.6600 in 2045 and 0.0921 in 2050. In the official series the Bolivian population almost stops growing, so the last years of the trajectory add very little. This is a property of the official projection and not a choice made here. The choice between the two candidate rules matters little in any case, since either gives 349,534 against 349,535 inhabitants in 2035 and 381,165 against 381,277 in 2050, a difference of 0.03 %.

The regional household size falls from 3.771 in 2024 to 3.434 in 2035 and 3.171 in 2050. Two other calibrations of the same functional form give 3.31 and 3.16 for 2050, which is reassuring about the calibration and says nothing about the form itself. The result also sits where a comparison would put it. The national census series gives 4.1 in 2001, 3.5 in 2012 and 3.1 in 2024, with 2.8 in rural areas [74], the observed range across Latin America runs from 3.0 in Chile to 4.9 in Guatemala [75], and no normative floor exists [76]. The trajectory is a scenario hypothesis calibrated on the study area and not an official projection, and a range of 2.9 to 3.6 in 2050 is carried forward as the main uncertainty of the demographic layer. Extending the official series past 2034 rests on assumptions that each have a known direction. Internal migration is assumed to slow in the same proportion as natural growth, which follows the method of the institute and pushes the result downwards. The ratio rule corrects the level of the official series and not its slope, so an error in the projected growth rate is inherited. The deceleration is abrupt at the end, the factor falling from 0.66 in 2045 to 0.09 in 2050, and no new trend reversal is assumed after 2034 although three municipalities already reverse inside the official series.

Table 17 gives the result. The population of the region gains 20.0 percent by 2050 and the number of households gains 42.7 percent. The two are not the same driver, and using the first as a proxy for the second misses half of the growth. Figure 11 shows the two paths against a common base.

Table 17: Projected population, households and household size. The cluster rows give households.

	2024	2035	2050	Growth to 2050
C1	10,933	12,942	14,873	+36.0 %
C2	802	998	1,204	+50.1 %
C3	40,298	47,958	55,641	+38.1 %
C4	16,612	20,801	25,487	+53.4 %
C5	15,564	19,084	22,994	+47.7 %
Households	84,209	101,783	120,199	+42.7 %
Population	317,517	349,534	381,165	+20.0 %
Household size	3.771	3.434	3.171	−15.9 %




Figure 11: Projected population and number of households in the region, indexed to 100 in 2024. The gap between the two curves is the fall in household size.

The mechanism is already visible in the observed data, without any projection. Between 2012 and 2024 all 21 municipalities gained households, and five of them lost inhabitants over the same period and gained households anyway, as Table 18 shows. Together these five hold 19.9 % of the households of the region.

The same mechanism would set the demand of the sufficiency bundle if a projection carried it, since the bundle is fixed per household: its total would follow the 42.7 percent of the households and not the 20.0 percent of the population.

Table 18: The five municipalities that lost population and gained households between the 2012 and 2024 censuses.

Municipality	Population	Households	Household size
Guayaramerín	−0.1 %	+17.6 %	4.34 to 3.68
Reyes	−12.6 %	+15.4 %	4.36 to 3.30
Bella Flor	−3.6 %	+30.8 %	3.76 to 2.77
Villa Nueva	−10.8 %	+20.2 %	4.73 to 3.51
Nueva Esperanza	−11.5 %	+4.9 %	3.81 to 3.21

4.10.3 The split between the three sources at each horizon

The dispersed share of each municipality comes from the spatial layer of Section 4.9. It is the Source C households of the municipality that live in a community classified as dispersed, divided by all its Source C households, and it is applied to the whole unconnected basin. It is zero in seven municipalities, Bolpebra, Cobija, Nueva Esperanza, Porvenir, Puerto Gonzalo Moreno, San Pedro and Villa Nueva, where the shared trunk reaches every community. It reaches 0.138 in Ingavi and 0.127 in Santos Mercado, the two most isolated municipalities of the region, and its regional mean weighted by the unconnected basin is 0.027. This is not the `share_dispersion` of Table 14, which is a ratio of energy rather than of households, and the two should not be read against each other.

Of the 9,325 Source B households of the census, 9,078 are matched to a community of the 2012 layer, and the standalone system is priced with the same technology catalogue as the optimisation.

Table 19 gives the regional split, and Table 45 in Appendix D gives the same by cluster.

Table 19: Households by source of electricity under the two access trajectories. Universal access means an empty Source C, not a Source A covering every household.

Trajectory	Year	Source A	Source B	Source C
Observed	2024	63,495	9,325	11,389
Access in 2035	2035	101,351	432	0
Access in 2035	2050	119,756	443	0
Access in 2050	2035	89,967	5,447	6,369
Access in 2050	2050	119,756	443	0

Both trajectories reach full access, both hold the same 120,199 households, and both leave the same 443 of them on standalone systems. Figure 12 shows where the two calendars separate. The slower one is trend based in its shape rather than in its total pace, since reaching full access in 2050 is a constraint of the scenario and a natural closure would not get there. What the calibration decides is how the effort is spread over the period, and the factor $\phi_m(2035)$ says how much falls in the first half, from 0.42 for Ingavi to 0.72 for Cobija.




Figure 12: Share of households with electricity from any source, by cluster. The two trajectories share the 2024 baseline and the 2050 endpoint, and differ only at the intermediate horizon.

The number of off grid households collapses. Under Access in 2035 it drops from 9,325 to 432 by 2035, then rises to 443 by 2050. Under Access in 2050 it falls from 9,325 to 5,447 and then to the same 443. Two mechanisms act on Source B, since line extensions pull households out of it and household growth pushes new ones in. The faster calendar finishes all the profitable extensions by 2035, so only the second mechanism is left after that year, and the number rises again by eleven households. The slower calendar is still building lines throughout, so it falls the whole way.

The same pattern appears cluster by cluster. C2 and C5 empty completely, because no community of Bolpebra or of Cobija falls outside the reach of the grid. C3 nearly empties too, from 2,274 to 81 households under Access in 2035, since it is the most urbanised cluster after C5. What remains in 2050 sits almost entirely in C1 and C4, with 146 and 207 households, which is 80 % of the regional total. In these two clusters a standalone system stays cheaper than a line at all three horizons.

The network is the yardstick of the calendar. Reaching every connectable community by 2035 means building 4,095 km of medium voltage line, about 372 km a year, and spreading the same programme to 2050 brings it down to 160 km a year. No national series of kilometres built per year exists for Bolivia. What can be cited is what ENDE reports about itself: nine distribution projects totalling 763.69 km of medium voltage line in 2021, across Cochabamba, Chuquisaca, Beni and Pando [77], then 178 km by its Beni subsidiary in 2024 and 98.72 km by its Oruro subsidiary in 2023 [78, 79]. None of these covers the cooperatives or the other distributors, so they are a floor rather than a total. The faster calendar therefore asks one region for about half of what ENDE built across four departments in its best documented year, for under three % of the Bolivian population. Figure 13 shows where that network runs and how little the two future horizons add to it, and Figure 34 in Appendix D gives the kilometres cluster by cluster. The network to build is decided by the settlements of today, since household growth adds 36 edges by 2035 and 9 more by 2050, on a tree that already holds 755.




Figure 13: Medium voltage tree at the 2050 household counts, each edge coloured by the horizon at which it first appears. The classification is run again at each horizon, so the three trees are not nested.

4.10.4 The projected demand of the connected households

Table 20 gives the projected Source A demand by end use and Table 21 the cluster totals that feed the demand files.

Table 20: Projected Source A demand by end use, in GWh/yr. The end use names are those of the demand file described in Section 4.8.3. The COOKING line carries only the electric cooking measured on the commercial tariff, and not the 111 GWh of household cooking, which is built separately from the census.

End use	2024	2035, access 2035	2035, access 2050	2050
FOOD_PRESERVATION	48.50	89.58	69.66	106.09
ELECTRICITY	54.51	84.22	74.51	99.43
SPACE_COOLING	32.59	52.29	44.58	61.93
LIGHTING_R_C	12.93	20.65	18.32	24.33
LIGHTING_P	9.71	15.09	13.52	17.79
MECH. ENERGY_IND	12.96	12.96	12.96	12.96
HEAT_LOW_T_HW	5.10	7.87	7.10	9.32
COOKING	2.31	3.47	3.17	4.12
MECH. ENERGY_COMM	1.48	2.23	2.04	2.65
Total	180.10	288.37	245.86	338.63

Source A demand is multiplied by 1.88 between 2024 and 2050, and the number of Source A households by 1.89. The intensity per connected household moves from 2,836 to 2,828 kilowatt hours, a change of three tenths of a %, so the growth of this demand is a phenomenon of access and demography rather than of consumption. The flatness is a coincidence rather than a law. Residential intensity does rise, by 9 %, as ownership follows access, while the industrial and residual sectors are frozen in level and are therefore diluted by the growing household count. The two happen to cancel at regional scale.

The same average hides very different clusters. In C2, where only 19 % of households were connected in 2024, the demand is multiplied by 11.3 while the connected households are multiplied by 7.7. In C5 the opposite happens, since demand rises by a factor of 1.33 against 1.54 for the households, so intensity falls by 14 %. The method builds this in, because ownership is tied to the access rate and a municipality already at 97.5 % has nothing left to gain.

Food preservation is the second end use of Source A in 2024, at 48.5 GWh behind the 54.5 GWh of the residual electricity line, and it is the one that grows fastest. It reaches 106.1 GWh in 2050 and takes first place, multiplied by 2.19 while the total is multiplied by 1.88, so its share of Source A demand rises from 27 to 31 %. The reason is direct, since refrigerator ownership is tied to the access rate and it is low in exactly the municipalities that gain the most access. This puts more weight on an assumption that Chapter 8 flags as the weakest of the model, the duty cycle of the refrigerator.

Two checks bound the method. Putting every appliance on the access relationship gives 340.36 GWh in 2050 and freezing every one of them gives 311.88, against 338.63 for the hybrid rule retained. The bracket is not symmetric. The upper bound sits 0.5 % above the hybrid rule and the lower one 7.9 % below it, so the hybrid rule is already close to the ceiling and the uncertainty on ownership can only lower the projected demand. Applying the whole machinery to the 2024 inputs returns the 2025 Source A demand to within 0.002 %, which confirms that the chain is continuous with the runs of Chapter 5. Three things do not move in this projection. The unit consumption of every appliance

Table 21: Projected Source A demand by cluster, in GWh/yr.

Cluster	2024	2035, access 2035	2035, access 2050	2050
C1	11.75	30.45	20.97	34.93
C2	0.24	2.24	1.29	2.70
C3	90.09	128.30	114.41	146.64
C4	22.88	65.89	48.66	81.14
C5	55.14	61.49	60.53	73.23
Total	180.10	288.37	245.86	338.63
Source A households	63,495	101,351	89,967	119,756
Intensity (kWh/household)	2,836	2,845	2,733	2,828

is held at its 2025 value, so no efficiency gain is credited and the 2050 demand comes out high. The industrial and residual tariffs are held at their 2025 level, which pushes the same demand down. Equipment ownership is also tied to the access rate, so a municipality already at full access keeps its 2024 equipment although that is where rising income pushes hardest. These biases work in opposite directions and none of them is corrected.

The external electricity available to C1 never binds. The 2025 runs leave it far above any import the cluster could make, and C1 draws 13.86 GWh over the year, which is the 11.593 GWh its households consume plus the network losses. The same rule is kept in projection, where the availability is set at roughly twice the total electricity demand of C1: 75 GWh per year in 2035 under Access in 2050, 86 GWh under Access in 2035, and 99 GWh in 2050. The projected Source A consumption of the C1 municipalities is a separate quantity, 18.56 GWh in 2035 under Access in 2050, 25.43 under Access in 2035, and 29.03 in 2050 under both.

4.10.5 The four projection scenarios

Everything not stated here is unchanged from the 2025 configuration: the clusters, the resource prices, the exchange topology, the technology catalogue and the time series. Table 22 lists the four scenarios, run over seven cases in total, three at 2035 and four at 2050.

Table 22: The projection scenarios. Diesel only carries the same demand as Access in 2050, so the pair isolates the supply rule. Access in 2035 and Access in 2050 share the supply rule, so the pair isolates the calendar.

Scenario	Access reached	Supply	Horizons
Diesel only	2050	existing plants, diesel only	2035, 2050
Access in 2050	2050	free investment	2035, 2050
Access in 2035	2035	free investment	2035, 2050
Access in 2035 with Brazilian import	2035	free investment, plus import	2050

The national commitment is more ambitious than either calendar, since the Bolivian Nationally De-

terminated Contribution targets universal access by 2030 [80]. That date falls before the first horizon of this work, so it is not modelled as a scenario, and Access in 2035 is the closest thing to it that the two horizons allow.

Diesel only is a bound rather than a forecast. Utility solar is already the cheapest way to produce a kilowatt hour in this region, as Chapter 5 shows, so a region that builds none of it by 2050 is not what anyone expects. Its purpose is to price the alternative that the two other trajectories reject, which is what keeping the present supply costs if nothing is built. Allowing it to build solar would not produce a third trajectory, it would produce Access in 2050.

The retirement of the 2025 fleet forces that scenario to be defined twice. In 2035 every generation and storage technology is locked to the capacity that exists, exactly as in the current system of Chapter 5. In 2050 the diesel floor has fallen to zero, so leaving the fleet locked would forbid the region to run any dispatchable plant at all. The diesel capacity is therefore released and the model sizes it freely, while utility solar and batteries stay locked at zero. Diesel only at 2050 is a region that keeps replacing its engines and builds nothing else. The Cobija solar plant is part of what it does not rebuild, 5.1 MW producing 7.8 GWh a year in 2035 and gone in 2050, which against a supply of some 370 GWh is about two percent served by diesel instead. The direction matters more than the size, since it pushes Diesel only away from the ones it is compared with.

The two trajectories start from the same installed capacity, that of the current system, which is the only fleet that physically exists in 2025. The brownfield rule of Equation 8 charges no capital on what already stands, so a trajectory anchored on a rebuilt system would inherit that fleet for free. The technology catalogue is identical in the two for the same reason, since Section 5.7 shows that blocking one technology moves the annual cost of a scenario by more than ten %.

The definition of the service also changes between the 2025 scenarios and these. There, the households of Sources B and C receive the normative sufficiency bundle. Here, a household that gains access becomes an ordinary Source A household and consumes at the rate the utility data give for connected households in its municipality. This second definition is empirical rather than normative, and it is the more realistic of the two once access means a line to the house. It means that this section does not continue the access transition scenario of Chapter 5, and that it answers a neighbouring question: what it costs to bring every household to the level of service that connected households in this region actually have. Table 23 gathers that difference and the five others that separate the two families of scenarios.

The grid extension is priced on the same basis as in 2025 and stays outside the generation cost. Reaching every connectable community costs 39.54 M€ of capital at the 2035 household counts and 40.68 M€ at the 2050 ones, which is 4.19 and 4.32 M€/yr once annualised over thirty years at the discount rate of the model. Both figures are for full access, so the second applies to the two trajectories and the first only to Access in 2035.

4.11 What this chapter establishes

The region is 84,209 households, of which 63,495 are supplied by the public service, 9,325 by their own panel or generator, and 11,389 by nothing at all. Three different methods estimate them, because the data available for each is different, and the heterogeneity between municipalities is what forces that separation: public grid coverage runs from 0.3 to 95.9 % inside one region.

Two demand definitions are built on the same territory and with the same model, the sufficiency bundle at 245.13 GWh of useful energy and the demand measured today. Neither is a forecast, and

Table 23: What differs between the five 2025 scenarios of Chapter 5 and the projection scenarios. Everything not listed is identical.

	2025 scenarios	Projection scenarios
Demand of a household that gains access	the normative sufficiency bundle	the observed consumption of a Source A household of the same municipality
Sufficiency demand	one scenario gives it to all 84,209 households	no scenario carries it
Capacity at the start of the run	the fleet that exists in 2025	the fleet of the previous horizon, chained through Equation 8
Investment costs	the 2025 catalogue	the same catalogue, with a declining ratio on utility solar, lithium storage, the household panel and the household battery
Reported cost	as returned by the optimisation	reconstructed at 2050, where the capacity inherited from 2035 is re-invoiced at its 2035 cost
External electricity available to C1	capped at the observed consumption	set high enough never to bind

the distance between them is what Chapter 5 measures. Cooking and hot water carry 74 % of that demand and neither has to be electric, which is why every cost of the next chapter is split between an electric perimeter and a non-electric one. The spatial layer supplies what the optimisation cannot see: 626 of the 697 communities are cheaper to reach with a line, 71 are cheaper to equip with standalone systems, and the lines cost 37.63 M€ of capital, added to the generation optimum wherever a scenario connects households that have no line today.

5 Results

This chapter reports what the five 2025 scenarios of Chapter 4 cost, what they emit, and how they supply the region. It follows the order in which the scenarios were defined. Two conventions on how a cost is measured come first, because every figure of the chapter depends on them. The current system comes next, since the four others are read against it, and each of those changes one thing at a time. A sensitivity on the water heating technology follows, then the comparison of the five, and a last section reports the seven runs of 2035 and 2050.

Rebuilding the supply saves more than any other single change, and no household changes what it consumes. Giving access to the households that have none costs more on a supply that has already been decarbonised than on the diesel system of today, which is not what one would expect. The emissions of the region are decided by whether the model may invest, not by how much electricity is consumed, so cost and emissions respond to two different decisions.

5.1 How the costs are reported

Every cost in this chapter follows the same two conventions, stated once here.

The first concerns the scope of the cost. The optimisation returns a single total annual cost that covers every energy service of the region, including cooking with wood and LPG and water heating with gas or with a solar collector. Dividing that total by an amount of electricity would mix a numerator that contains non-electric fuels with a denominator that does not. The cost is therefore split in two. The electric scope gathers every technology that produces or consumes electricity, the batteries, the network, and the electricity and diesel consumed by those technologies. The non-electric scope gathers non-electric cooking, gas and solar water heating, and thermal storage.

The second concerns the denominator. In EnergyScope each end use has its own layer, so the electricity drawn by a refrigerator or a light bulb does not appear on the ELECTRICITY line of the demand file. The denominator used here is the sum of two terms: the electricity consumed by the end-use technologies, and the demand served directly from the electricity layer. Network losses, storage cycling and inter-cluster exchanges are excluded, since none of them is an end use.

The unit cost reported in the sections below is the supply cost. It covers the generation, the storage and the network, which is what compares to a tariff or to a levelised cost of electricity. The end-use appliances are outside it, because the technology database gives no investment and no maintenance cost to the refrigerator, which is the largest single electricity consumer in every scenario. The same database gives a construction footprint of zero to the network technologies, so the 3,948 km of medium voltage line, the inter-cluster lines and the hot water tanks carry no embodied carbon in any emission figure of this chapter. Both gaps understate the scenarios that build the most infrastructure.

5.2 Current system

The current system scenario is a costed picture of the present situation. The demand is the measured one, the existing plants are the only ones available, and the model can do nothing but dispatch them. It is the reference against which the other four scenarios are read.

The system supplies 209.6 GWh of electricity over the year, of which 176.6 GWh reaches an end use. The difference is the network loss of Section 4.8.1. No storage is installed in this scenario, so nothing is lost to cycling either. Diesel gensets produce 184.9 GWh, or 88 % of the total. The national grid supplies 13.9 GWh to C1, the only cluster connected to it, and the Cobija solar plant produces 7.8 GWh. Home systems produce 3.0 GWh, which is 1.4 % of the regional supply and the entire consumption of the 9,325 off-grid households.

The region therefore runs on diesel. The gensets and the household generators burn 508 GWh of fuel to deliver 187.5 GWh of electricity, an average conversion efficiency of 37 %. Diesel machinery adds 3.8 GWh, which brings the diesel bill of the region to about 512 GWh a year. That efficiency is in the usual range for diesel gensets. The size of the fuel bill comes from the price of diesel and from its share of the supply, not from the efficiency of the engines.

Figure 14 traces the annual flows of that system from resources to services. Two chains carry the region and they barely touch: diesel into electricity, and LPG into cooking.




Figure 14: Annual energy flows of the current system in 2025, in GWh, from resources to end uses. Network losses appear as an end use of their own.

The total annual cost is 66.51 M€/yr and it divides in three. Supplying the electricity costs 58.21 M€/yr, which covers the plants, the fuel they burn, the import and the network. The electrical equipment of the end uses adds 2.45 M€/yr, dominated by air conditioning at 1.97 M€/yr. The remaining 5.85 M€/yr falls outside the electricity scope altogether, mostly cooking with LPG. Only the first of the three is a supply cost, following the convention of Section 5.1.

The structure is almost entirely operational: 62.80 M€/yr of resource cost against 2.08 M€/yr of investment and 1.64 M€/yr of maintenance. The investment that does appear is not generation. The plants are brownfield and carry no capex, so what is left is the end-use equipment the model has to build to serve the demand.

Diesel alone accounts for 56.55 M€/yr, which is 97 % of the supply cost. Divided by the 187.5 GWh it produces, the fuel represents 0.30 €/kWh at the point of generation, before anything else is added.

Maintenance and the network bring the supply cost to 0.33 €/kWh of electricity delivered to an end use.

Table 24 gives the cost and the emissions by cluster.

Table 24: Annual cost and emissions by cluster, current system.

Cluster	Cost (M€/yr)				Emissions (ktCO ₂ eq/yr)	
	Investment	Maintenance	Resources	Total	Operation	Construction
C1	0.17	0.09	1.59	1.86	9.83	0.17
C2	0.01	0.00	0.08	0.09	0.57	0.02
C3	0.89	0.78	34.06	35.74	113.97	1.37
C4	0.37	0.20	9.17	9.74	33.07	0.43
C5	0.64	0.56	17.90	19.09	57.25	1.00
Total	2.08	1.64	62.80	66.51	214.69	2.98

C3 and C5 together carry 82 % of the regional cost. Both host the large isolated diesel systems of the region, Riberalta and Guayaramerín on one side and Cobija on the other. C1 is the cheapest cluster per unit served, because it is the only one that draws on the national grid instead of burning fuel locally. C2 carries 0.1 % of the regional cost. Bolpebra is one small municipality, and the clustering keeps it apart for its economic structure and not for its energy volume.

Operational emissions are 214.7 ktCO₂eq per year and come almost entirely from diesel combustion. Construction emissions add 2.98 ktCO₂eq per year, for a total of 217.7 ktCO₂eq. This scenario builds almost nothing. The model charges construction emissions on the total installed capacity while it charges investment cost only on the capacity built beyond what already exists, so the existing diesel fleet carries its embodied emissions without carrying its capital cost. The 2.98 ktCO₂eq is therefore not the footprint of anything built in 2025.

5.3 Supply transition

The supply transition scenario keeps the demand of the current system unchanged and lets the model invest freely. No household changes what it consumes. The only thing that changes is what the region is allowed to build. The question it answers is narrow on purpose: how much of the present cost is caused by the way electricity is produced, rather than by how much of it is consumed.

The optimiser installs about 154 MW of utility PV, of which 149 MW is new, concentrated in C3 (82 MW), C5 (51 MW) and C4 (20 MW). It pairs this with 225 MWh of lithium-ion storage in the same three clusters: 122 MWh in C3, 73 MWh in C5 and 29 MWh in C4. C1 receives none, because it has no local generation to firm and draws on the national grid instead.

These two technologies carry the whole of the new investment. Utility PV costs 17.62 M€/yr in investment and maintenance, and the batteries 12.76 M€/yr. Together they are 89 % of the annual capital charge of the scenario, and everything else, from the lighting to the stoves to the inter-cluster lines, shares the remaining tenth. The existing diesel fleet stays installed but barely runs. Its output falls from 184.9 to 9.3 GWh, and regional diesel consumption falls from about 512 to 34 GWh, a reduction of 93 %.

Utility PV supplies 220.6 GWh, or 89 % of the 246.6 GWh delivered by the system. The national grid contributes 5.6 %, the remaining diesel 3.8 %, and home systems 1.2 %. In one step the region goes from a system that is 88 % diesel to one that is 89 % solar, without a single household changing its behaviour.

The mix is maintained hour by hour through two mechanisms. The batteries move the midday production into the evening and the night, and the diesel fleet remains in place for the hours when neither the sun nor the storage capacity is available, which is what its remaining 9.3 GWh serve. The energy balance therefore closes at every hour of the 16 typical days. However, achieving an energy balance is not the same as operating the system. In a grid where 89 % of the energy arrives via inverters, little rotating inertia remains, so the frequency control that the diesel engines provide today must come from elsewhere. This question sits outside the model, under the assumption of Section 3.3.1, and the tools referenced there can answer it.

The electricity delivered to end uses rises from 176.6 to 199.8 GWh even though the useful demand is identical in the two scenarios. The 23.2 GWh of difference comes from end uses that change carrier, electric cooking above all. LPG consumption falls accordingly, from 4.46 to 3.67 M€/yr.

Nothing in the scenario asked for this. The cooking mix is left entirely free, with no minimum or maximum share imposed on any stove, so the switch is the optimiser choosing the cheaper option once solar electricity is available. Building solar generation therefore does not only displace diesel. It also pulls onto the electricity layer thermal demand that subsidised fuels served before.

The total annual cost falls to 42.15 M€/yr, a reduction of 37 % against the current system. The structure of the cost turns over completely. Resource cost drops from 62.80 to 8.32 M€/yr, while investment and maintenance rise from 3.71 to 33.83 M€/yr. The system stops paying for fuel every year and starts paying back capital instead. The supply cost of electricity falls from 0.33 to 0.18 €/kWh.

Table 25: Annual cost and emissions by cluster, supply transition.

Cluster	Cost (M€/yr)				Emissions (ktCO ₂ eq/yr)	
	Investment	Maintenance	Resources	Total	Operation	Construction
C1	0.17	0.09	1.59	1.86	9.83	0.17
C2	0.17	0.03	0.06	0.26	0.38	0.10
C3	15.27	2.81	3.69	21.77	24.83	7.52
C4	3.83	0.70	1.34	5.87	9.39	1.93
C5	9.05	1.71	1.64	12.40	10.37	4.51
Total	28.49	5.34	8.32	42.15	54.80	14.24

Table 25 gives the cost and the emissions by cluster, and C1 is identical to the current system to the fourth decimal. It has no local diesel plant and no local generation to replace, so it stays on imported grid electricity whatever the investment rules are. Every euro of the saving is made in the four other clusters, and more than half of it in C3 alone, which goes from 34.06 M€/yr of fuel to 15.27 M€/yr of investment.

Operational emissions fall from 214.7 to 54.8 ktCO₂eq per year. Construction emissions rise from 2.98 to 14.24, which is the embodied footprint of the new solar and battery fleet, annualised over its

lifetime. The total falls from 217.7 to 69.0 ktCO₂eq, a reduction of 68 percent.

5.4 Universal access

This is the scenario a planner faces. What does it cost to bring a decent level of service to the 20,714 households that are badly served or not served at all, without changing anything for those who already have electricity. The grid-connected demand stays at its measured level, and the households of Sources B and C move to the sufficiency bundle. The production of the existing plants is frozen to what they deliver in the current system, using the mechanism of Section 3.4.

The existing diesel fleet keeps producing exactly what it produces today: 184.9 GWh from utility gensets and, for the households that stay dispersed, a bounded output from home generators. The new demand adds about 26.8 GWh of electricity to end uses, raising the regional denominator from 176.6 to 203.4 GWh. The model builds the generation needed to cover this increment, plus the solar water heaters for the hot water of the new households.

The 20,714 households that move to the sufficiency bundle are not spread evenly. C4 holds 7,963 of them, which is 48 % of its own population: nearly one household in two in rural Pando has no public supply today. C3 holds 6,372, but they represent only 16 % of a cluster that is already 84 % connected. C1 holds 5,090, or 47 % of its households. C2 is the least served of all, with only 19 % of its 802 households on the public grid. C5 is the opposite case, with 96 % of Cobija already connected and only 643 households left out.

This distribution explains the shape of the results that follow. The cost of the access is concentrated in C1 and C4, where the households are many and the communities are far apart, while the cost of the system as a whole stays concentrated in C3 and C5, where the diesel is burnt.

The new generation is modest. Utility PV reaches 34.5 GWh. This is a fraction of the 220.6 GWh built in the supply transition, because the model cannot re-optimize the existing supply and only adds what the new households need. The installed PV capacity is about 22 MW, compared with 154 MW in the supply transition. No new lithium-ion battery is deployed at cluster scale.

Home systems serve the households that remain dispersed after the spatial classification. Their production is 0.54 GWh, in C1, C3 and C4 alone, since the classification leaves no dispersed community in C2 or in C5. Table 26 compares their output with the target demand.

Table 26: Home system production against the target demand of dispersed households, universal access.

	C1	C2	C3	C4	C5
Target (GWh/yr)	0.179	0	0.111	0.242	0
Home systems (GWh/yr)	0.181	0	0.113	0.246	0
Deviation	+0.8 %	—	+2.1 %	+1.5 %	—

The home systems deliver the target demand of the dispersed households to within 2.1 % in every cluster. The residual comes from the shape of the dispersion constraint, which is enforced hour by hour on a value calibrated on annual totals, as Section 4.9.3 describes. C2 and C5 hold no dispersed community, so their home system technologies are bounded to zero.

The solar water heater is available here and not in the current system, where the connected households do not own one. It is opened for the new households of Sources B and C.

The hot water demand is 22.73 GWh and the heat layer produces 23.8 GWh to serve it, the difference being the tank losses again. Of that production the solar collector serves 17.2 GWh (72 %), the existing gas water heaters of Source A contribute 2.4 GWh (10 %), and the electric resistance serves 4.2 GWh (18 %). The losses are proportionally smaller than in universal sufficiency, at 4.5 percent of the production against 5.5, because the two electric and gas paths can produce at the hour the water is drawn and skip the tank. The gas contribution is identical to the current system, down to the third decimal, because the production of the existing gas heaters is frozen by the same mechanism as the plants. The water heating of Source A is therefore held fixed, and the solar collectors serve the new households only.

Part of the cost difference against the current system therefore comes from a technology becoming available, and not only from the access itself. Section 5.7 separates the two.

The generation cost is 69.02 M€/yr, which is 2.51 M€/yr above the current system. Adding the grid extension brings the total to 73.01 M€/yr, or 6.50 M€/yr above. Per new household this is 314 €/yr, or about 26 € per month. These are annualised costs that include the capital recovery of the lines and of the generation equipment. They are not an upfront expenditure, but they are the yearly budget the region would have to find.

Table 27: Annual cost and emissions by cluster, universal access (generation only).

Cluster	Cost (M€/yr)				Emissions (ktCO ₂ eq/yr)	
	Investment	Maintenance	Resources	Total	Operation	Construction
C1	0.40	0.12	1.91	2.42	11.05	0.14
C2	0.04	0.01	0.05	0.09	0.48	0.01
C3	1.85	0.98	33.68	36.50	112.87	1.83
C4	0.76	0.24	8.80	9.80	32.04	0.35
C5	1.65	0.82	17.74	20.20	56.80	1.72
Total	4.69	2.16	62.17	69.02	213.24	4.05

Table 27 gives the cost and the emissions by cluster. The supply cost of electricity falls slightly, from 0.33 to 0.30 €/kWh without the grid extension, because the new PV in C3 and C5 serves the added demand more cheaply than the diesel average. With the grid extension it rises to 0.32 €/kWh, still below the current system. Figure 15 compares the two scenarios side by side.




Figure 15: Current system against universal access: production, cost and emissions by cluster. The increment is concentrated where the new households are, and the existing diesel production is identical in the two scenarios.

Unlike the supply transition, this scenario does not flip the supply mix. Diesel still produces 184.9 GWh, exactly as in the current system. The PV that is built covers the increment, not the base. This is by design. The frozen production of the existing plants prevents any substitution, so the cost comparison measures the access alone, without the saving that would come from switching the existing fleet to solar.

Changing the access and the supply at the same time gives a different picture, and it is what universal sufficiency does in Section 5.6. Section 5.8 splits the gap between the two into three steps, and shows that the cost of the access depends on whether the supply has already been rebuilt.

Total emissions are 217.3 ktCO₂eq, marginally below the current system (217.7). Operational emissions are lower (213.2 against 214.7), because part of the household diesel generation that served Source B has been replaced by home PV. The small increase in construction emissions reflects the embodied carbon of the new panels and collectors.

The access is therefore carbon neutral to first order. The households that gain access are a quarter of the total, but their demand is only 15 % of the regional electricity. The existing diesel fleet sets the emissions of the region, and the new demand barely moves them.

Figure 16 decomposes the 6.50 M€/yr increment. The generation increment is 2.51 M€/yr: 2.04 M€/yr of new utility PV, 0.94 M€/yr of solar water heaters and their tanks, 0.48 M€/yr of additional grid import, 0.09 M€/yr of refrigeration and lighting equipment and 0.05 M€/yr of home systems, against which 1.11 M€/yr of avoided diesel is set. The grid extension adds 3.99 M€/yr. Almost two

thirds of the cost of access is therefore the lines, not the generation.




Figure 16: Decomposition of the annual cost of universal access, from the current system (66.51 M€/yr) to the total including grid extension (73.01 M€/yr). The extension dominates the increment.

One of these components carries a technology choice that the current system does not have. DEC_SOLAR is blocked in the current system and opened here. Without it, the generation increment would rise by 0.47 M€/yr (Section 5.7), or about 7 % of the total access cost. The choice of hot water technology is a real part of the answer, but not a dominant one. In universal sufficiency, where the collector serves four times as much demand, the effect is much larger.

5.5 Access transition

This scenario has the demand of universal access and the investment rules of the supply transition, as defined in Section 4.3. The 20,714 households of Sources B and C receive the sufficiency bundle, the grid-connected households keep their measured demand, and nothing is held fixed.

The generation cost is 47.07 M€/yr. The grid extension of Section 4.9 adds 3.99 M€/yr, which brings the total to 51.06 M€/yr, or 606 € per household of the region. The supply cost of electricity is 0.18 €/kWh, the same value as the supply transition, since the two scenarios build under the same rules.

End uses receive 219.2 GWh, above universal access and above the supply transition. The grid-connected households keep a measured demand several times larger than the bundle, and the new households add the bundle on top of it. Part of the increase also comes from cooking moving onto the electricity layer, the same effect described in Section 5.3.

Diesel burnt in the region is 40 GWh, against 34 in the supply transition, and emissions rise from 69.0 to 73.6 ktCO₂eq per year. The reason is that the cheapest way to serve the hours the solar cannot cover is to restart the gensets the transition had switched off. Their output rises from 9.3 to 13.4

GWh. Section 5.8 uses this scenario to measure what that costs, and finds that giving access on a decarbonised system is more expensive than giving the same access on the diesel system it replaced.

5.6 Universal sufficiency

In this scenario every household of the region receives the same normative service bundle, whether it is connected today or not. The model is free to build anything, and the sufficiency demand replaces the measured one for all 84,209 households. The scenario answers a question that no planner can avoid: what would it cost if every household in the region consumed exactly what the sufficiency framework defines as a minimum acceptable level of service.

The system supplies 156.3 GWh of electricity, of which 126.2 GWh reaches an end use. This is less than the 209.6 GWh of the current system, not because fewer people are served but because the sufficiency bundle consumes less per household than the measured demand of the grid-connected population.

Utility PV produces 142.2 GWh, or 91 % of the supply. The national grid provides 11.5 GWh to C1, the only cluster with physical access to it. Diesel gensets produce 2.0 GWh, almost all of it in C3 and C5, and home systems another 0.6 GWh. The region is therefore 91 % solar in this scenario, compared with 88 % diesel in the current system. Figure 18 shows the supply by cluster.

The optimiser builds about 95 MW of utility PV, of which 90 MW is new, concentrated in C3 (51 MW), C4 (20 MW) and C5 (20 MW). It pairs this with 147 MWh of lithium-ion storage, 78 MWh of it in C3, 33 MWh in C4 and 30 MWh in C5. The storage sits where the solar is, since almost no dispatchable plant is left to cover the evening. This is also the scenario where the operating questions of Section 5.3 weigh the most, since the diesel fleet runs at 2.0 GWh against 9.3 there and almost no dispatchable plant remains. The model closes the energy balance, and the frequency and reserve of such a system remain for the tools of Section 3.3.1.

Figure 17 shows how that system runs over four typical days in C3. Midday solar covers the demand several times over, and the surplus charges the batteries that carry the evening and the night. On the one day of the four where the solar resource is weak, the batteries do not reach the evening and the gensets run instead, in the morning and again after sunset.




Figure 17: Hourly electricity dispatch in C3 over four typical days, universal sufficiency 2025. Production sits above the axis and the sinks that are not an end use below it, so the algebraic sum of the bars equals the demand line at every hour.

C2 barely appears. As a single municipality of 802 households it is too small to justify local PV at scale, and it draws most of its electricity from C5, 15 km away, through the existing inter-cluster line.




Figure 18: Electricity supply by cluster and technology, universal sufficiency 2025. Only C1 imports from the national grid; the imports of C2 and C4 come from neighbouring clusters.

The solar water heater is available in this scenario, and its effect is large. The region needs 71.67 GWh of hot water and the decentralised heat layer produces 75.8 GWh to deliver it, the 4.2 GWh of difference being the self-discharge of the daily tanks. Of that production the solar collector serves 72.5 GWh (96 %), the electric resistance covers 3.4 GWh (4 %) as backup, and the gas boiler is negligible. The tank losses are high here because almost all the heat is made in daylight and drawn in the evening, so nearly everything transits through storage. Figure 19 shows the split by cluster.

The collector fleet also spills what it cannot place. Over the year it curtails 17.94 GWh of heat, a fifth of what it could have produced, and more than half of that in C3 alone. The tank is a daily one and the demand is spread evenly over the year, so a collector sized for the hours when the resource is weak has nowhere to put its output when the resource is strong. Chapter 8 returns to the flat hot water profile that produces this effect.

The installed collector capacity is about 48 MW_{th} across the five clusters, at a cost of 4.32 M€/yr in investment and maintenance. This is the third largest cost item of the scenario, behind utility PV and batteries. The alternative is expensive. Without the collector, the 72.5 GWh of heat would have to come from electric resistance, and every kilowatt hour of heat would become a kilowatt hour of electricity that the PV and battery fleet must supply.

Section 5.7 measures what that alternative would cost.




Figure 19: Water heating by technology and cluster, universal sufficiency. The solar collector serves 96 % of the regional demand; the electric resistance covers the rest.

The generation cost is 31.28 M€/yr. This is less than half of the current system (66.51 M€/yr), even though it serves 20,714 more households. The non-electric scope accounts for 9.28 M€/yr, or 30 % of the total, dominated by the solar water heaters and by LPG for cooking. This share is much higher than in the current system (9 %), because the electricity cost has dropped while the cooking and hot water costs have not.

The supply cost of electricity is 0.17 €/kWh, and the end-use appliances add less than half a cent to it, since the technology database charges almost nothing for them. Both figures are far below the supply cost of the current system (0.33 €/kWh).

Table 28 gives the cost and emissions by cluster.

Table 28: Annual cost and emissions by cluster, universal sufficiency (generation only).

Cluster	Cost (M€/yr)				Emissions (ktCO ₂ eq/yr)	
	Investment	Maintenance	Resources	Total	Operation	Construction
C1	1.15	0.18	1.35	2.68	8.73	0.40
C2	0.25	0.04	0.04	0.33	0.38	0.11
C3	11.49	2.01	2.34	15.84	20.91	5.75
C4	4.88	0.75	0.95	6.58	9.12	2.07
C5	3.97	0.83	1.05	5.85	8.40	2.37
Total	21.74	3.80	5.74	31.28	47.53	10.69

The cost structure is dominated by investment (69 %), not by fuel (18 %). The three largest items are utility PV at 10.72 M€/yr, lithium-ion batteries at 8.33 M€/yr and solar water heaters at 4.32 M€/yr. Together they carry 75 % of the annual cost. LPG for cooking is the largest resource cost at 3.68 M€/yr, ahead of diesel at 1.37 M€/yr. Figure 20 shows the decomposition.




Figure 20: Cost decomposition by cluster (left) and top cost items across the region (right), universal sufficiency 2025.

The generation cost assumes the grid reaches every household. It does not. Section 4.9 classifies 71 of the 697 communities as dispersed, and the extension to the other 626 costs 37.63 M€ of capital, or 3.99 M€/yr annualised. The total cost of universal sufficiency including the network is therefore 35.27 M€/yr.

Total emissions are 58.2 ktCO₂eq per year, of which 47.5 are operational and 10.7 come from the construction of the new infrastructure. This is 73 % below the current system (217.7 ktCO₂eq). Operational emissions come mainly from cooking and from the residual diesel generation in C3 and C5. The 10.7 ktCO₂eq of construction is the annualised footprint of the new fleet, charged every year of its lifetime.

The emission database gives a zero construction footprint to the network technologies, so the 3,948 km of new medium-voltage lines carry no embodied carbon in these figures. Chapter 8 reports it.

A system that serves the sufficiency bundle to the entire region costs 35.27 M€/yr, which is 47 % less than what the current system pays to serve a smaller and unequal demand. Two things change at the same time, the demand and the supply. The result also depends on one technology that is not standard in the region, the solar water heater, and Section 5.7 measures how much.

5.7 The hot water technology lever

One technology carries an unusual weight in universal access and in universal sufficiency. The solar thermal collector serves 96 % of the domestic hot water in universal sufficiency and 72 % of it in universal access. It is also the technology whose availability was decided in this work rather than inherited. It exists in the national catalogue but was never activated in the national runs, and it is blocked in the current system and in the supply transition because the households of Source A do not own one today.

A result that rests on a choice like this has to be tested. Two additional runs were produced on the exact final configuration of universal access and universal sufficiency, with the solar collector blocked and nothing else changed. The hot water demand does not move, and only the technology allowed to serve it does, as Table 29 shows.

Table 29: Effect of blocking the solar water heater. Reference runs against sensitivity runs.

	Universal access		Universal sufficiency	
	Reference	Blocked	Reference	Blocked
Generation cost (M€/yr)	69.02	69.49	31.28	34.58
Difference		+0.47		+3.30
Difference (%)		+0.7		+10.6
Solar collector (GWh/yr)	17.2	0	72.5	0
Electric resistance (GWh/yr)	4.2	20.3	3.4	76.0
Electricity consumed by end-use technologies (GWh/yr)	147.4	163.5	120.1	196.8
Emissions (ktCO ₂ eq/yr)	217.3	218.8	58.2	62.0

Blocking the collector raises the cost of universal access by 0.47 M€/yr, which is under one %. It raises the cost of universal sufficiency by 3.30 M€/yr, which is nearly eleven %. The difference is entirely a matter of volume. The collector serves 17.2 GWh of heat in one scenario and 72.5 GWh in the other, because universal access only equips the households that gain access while universal sufficiency equips every household of the region.

To put the 3.30 M€/yr in perspective, the grid extension that connects 626 communities costs 3.99 M€/yr. In universal sufficiency, choosing between a solar collector and an electric resistance on the roof of every house is worth four fifths of the entire medium voltage network. It is a decision of the same order as whether to build the lines at all.

Heat produced by the sun never touches the electricity layer. Heat produced by a resistance does, one kilowatt hour for one kilowatt hour. Blocking the collector therefore moves 72.5 GWh of demand from the heat layer onto the electricity layer in universal sufficiency, and the electricity consumed by the end-use technologies rises from 120.1 to 196.8 GWh, an increase of 64 %. This figure counts only what the appliances draw, and it is not the denominator used for the unit costs of Section 5.1, which also includes the demand served directly from the electricity layer.

That new demand has to be produced. The annualised investment rises by 1.77 M€/yr and the main-

tenance by 1.12 M€/yr, almost entirely in additional solar panels. The fuel cost moves little, since the extra diesel is nearly offset by the LPG the system stops burning. What grows is the size of the system rather than the fuel it burns. Of the 76.8 GWh of additional electricity that the end-use technologies consume, 72.7 go to the electric resistance and 4.1 to cooking, which the optimiser moves a little further from LPG to electricity. Figure 21 shows the three effects together.




Figure 21: Effect of blocking the solar water heater in universal access and universal sufficiency: annual cost, hot water supply by technology, and electricity consumed by the end-use technologies.

Total emissions rise from 58.2 to 62.0 ktCO₂eq in universal sufficiency, an increase of 6.5 %. Operational emissions barely move, from 47.5 to 47.3. Construction emissions rise from 10.7 to 14.6, an increase of 37 %, which is the embodied carbon of the extra panels and batteries. A design that avoids building infrastructure is therefore cleaner than one that builds more of it, even when the extra

infrastructure is renewable.

No unit cost is reported for the two sensitivity runs, since their denominators differ from those of the reference runs.

5.8 Comparing the five 2025 scenarios

The five scenarios differ along two axes. One is the supply, whether the model may build new generation or only run what already exists. The other is the demand, how much electricity each part of the population is assumed to need. Table 30 puts them side by side and Figure 24 gives the same comparison graphically.

Table 30: The five 2025 scenarios side by side. All costs are annual. The supply cost excludes the grid extension in every column.

	Current system	Supply transition	Universal access	Access transition	Universal sufficiency
Households at sufficiency level	0	0	20,714	20,714	84,209
Generation cost (M€/yr)	66.51	42.15	69.02	47.07	31.28
Grid extension (M€/yr)	—	—	3.99	3.99	3.99
Total cost (M€/yr)	66.51	42.15	73.01	51.06	35.27
Cost per household (€/yr)	790	500	867	606	419
Electricity to end uses (GWh/yr)	176.6	199.8	203.4	219.2	126.2
Supply cost (€/kWh)	0.33	0.18	0.30	0.18	0.17
Diesel burned, all uses (GWh/yr)	512	34	502	40	12
Emissions (ktCO ₂ eq/yr)	217.7	69.0	217.3	73.6	58.2

The cost per household divides the total by the 84,209 households of the region in all five cases, so that the five figures compare. Cooking is served for every household in every scenario, which is what makes this denominator the right one. The supply cost excludes the grid extension everywhere, following the convention of Section 5.1, since a line is an infrastructure cost and not a cost of producing a kilowatt hour. Adding the extension back raises it to 0.32 €/kWh for universal access, and to 0.20 for the access transition and for universal sufficiency.

Figure 22 shows what produces that electricity in each scenario. The investment rule turns the mix over, from 88 % diesel to about 90 % solar, while the demand definition only changes how much of it the region needs.




Figure 22: Annual electricity production by technology in the five 2025 scenarios, in GWh. The two scenarios that may not invest keep the diesel fleet at its present output.

The generation cost falls by 35.23 M€/yr between the current system and universal sufficiency. Table 31 splits that gap into three steps. Each one changes a single thing and leaves the rest alone. The three add up to the total exactly, because they are consecutive and share their endpoints, so there is no residual left to attribute.

Table 31: The gap between the current system and universal sufficiency, decomposed into three consecutive steps. Generation cost only.

Step	What changes	M€/yr
Current system to supply transition	the supply decarbonises	−24.36
Supply transition to access transition	Sources B and C reach the bundle	+4.92
Access transition to universal sufficiency	Source A comes down to the bundle	−15.79
Total	66.51 to 31.28	−35.23

Figure 23 draws that decomposition. The two savings dominate it, and the cost of the access is small against either of them.




Figure 23: The three steps between the current system and universal sufficiency, generation cost only. Each bar changes one thing and leaves the rest alone, and the three sum exactly to the 35.23 M€/yr gap.

The first step rebuilds the supply and touches nothing else. No household changes what it consumes. It saves 24.36 M€/yr, which is 37 % of the annual cost of the current system, and more than the two other steps combined in absolute value. Section 5.3 describes what the model builds to get there.

Access costs more on a clean system than on a dirty one. The second step gives the sufficiency bundle to the 20,714 households of Sources B and C. It costs 4.92 M€/yr of generation. The same step measured against the current system, which is the universal access scenario of Section 5.4, costs 2.51 M€/yr. Same households, same bundle, same lines, and 2.42 M€/yr of difference between the two.

The reason is what the existing diesel fleet is doing in each case. In the current system it runs at 184.9 GWh and carries no capital cost, since the brownfield rule of Equation 8 charges capital only on capacity built beyond what already exists. It covers the night and the peak for nothing. The access increment therefore only has to buy energy, and 2.04 M€/yr of new utility solar is enough. It even displaces a little household diesel, so the fuel bill falls by 1.11 M€/yr.

In the supply transition the same fleet is still installed but almost idle, at 9.3 GWh. When the new demand arrives, the cheapest way to serve the hours that solar cannot reach is to run those engines again. Their output rises to 13.4 GWh, and the diesel burnt in the region rises from 34 to 40 GWh. The model also adds 16 MWh of lithium-ion storage, where the same step on the current system adds none at all.

The emissions do not follow the same logic. Access on the current system leaves them where they are, 217.7 against 217.3 ktCO₂eq per year, since the home diesel it displaces offsets what the new demand adds. Access on the decarbonised system adds 4.6, from 69.0 to 73.6, and most of that increase is operational rather than embodied. Connecting households to a clean system is not carbon free, because part of the marginal supply comes back from the engines the transition had switched

off.

The third step brings Source A down to the same bundle as everyone else. It saves 15.79 M€/yr, which is more than three times the cost of the access and almost four times the grid extension. Sufficiency is not only a floor under the households that have too little. It is also a ceiling over those that consume more than the bundle, Cobija in particular, and the ceiling is worth far more than the floor costs. Section 5.6 gives the resulting system.

One small effect is folded into this step. Universal sufficiency opens six inter-cluster exchange pairs where the other scenarios open three, so the exchange topology is not identical on both sides of the comparison. The entire HVAC line cost of the access transition is 0.24 M€/yr, so the difference in topology can be worth at most a few hundredths of a million against the 15.79. It is reported rather than corrected.

Figure 25 places the five scenarios on a cost against emissions plane, and the emissions axis splits them into two groups with nothing in between. The two scenarios in which the existing plants cannot be replaced sit at 217 ktCO₂eq per year. The three in which the model may invest freely sit between 58 and 74. Emissions are decided by one thing only, whether the supply may be rebuilt, and the demand assumption barely moves them.

Cost behaves differently. It varies inside each group, and there it is the demand that drives it: 66.51 against 73.01 M€/yr in the diesel group, and 35.27 against 51.06 in the solar group. There is no trade off between affordability and decarbonisation in this region, but there is no coupling either. The two objectives respond to two different decisions.

Universal access is the clearest illustration. It sits above the current system in cost and level with it in emissions. That is what the scenario was built to show. Connecting the households that have no electricity does not clean up the system, because the system is dirty for a reason that has nothing to do with them. The 20,714 new households add 15 % to the regional electricity demand, while the existing diesel fleet that produces 88 % of the electricity stays exactly as it is.




Figure 24: The five 2025 scenarios compared: annual cost split between generation and grid extension, electricity delivered to end uses, and greenhouse gas emissions split between operation and construction.




Figure 25: Annual cost against greenhouse gas emissions for the five 2025 scenarios. Marker size is proportional to the number of households served at the sufficiency level.

C3 and C5 carry between 69 and 82 % of the regional cost in every scenario. They host the two large isolated diesel systems, Riberalta with Guayaramerín on one side and Cobija on the other. Every euro of the supply saving is made there. C1 is the exception. It draws on the national grid, has no local plant to replace, and its cost is identical to the fourth decimal between the current system and the supply transition.

The five scenarios do not deliver the same basket of services, so their unit costs are not fully comparable. Cooking is electrified in the supply transition (20.0 GWh), in the access transition (19.5 GWh) and in universal sufficiency (17.7 GWh), and not at all in the current system or in universal access. A scenario that moves a service onto the electricity layer raises its own denominator without serving anyone better. The unit costs of Table 30 should therefore be read as an order of magnitude, not as a like for like ranking. The total annual cost, which covers every energy service in every scenario, is the safer comparison.

5.9 Results for 2035 and 2050

This section reports the seven runs defined in Section 4.10.5. One difference with the five scenarios above governs the whole reading: a household that gains access here becomes an ordinary connected household and consumes at the observed rate, instead of receiving the normative sufficiency bundle. Two pairs carry the reading. Diesel only and Access in 2050 serve the same households at each horizon, so the pair gives the price of the supply rule. Access in 2035 and Access in 2050 follow the same supply rule, so the pair gives the price of the calendar. All costs are reconstructed as defined in

Section 3.6.4, and none of them includes the grid extension. Table 32 gives the seven runs.

Table 32: The seven projection runs. The emissions of the import scenario use the central emission factor discussed below.

Horizon	Scenario	Cost (M€/yr)	Emissions (ktCO ₂ eq/yr)	Diesel (GWh/yr)
2035	Diesel only	89.45	286.8	689
2035	Access in 2050	44.06	76.7	20
2035	Access in 2035	49.29	81.1	23
2050	Diesel only	123.41	380.7	945
2050	Access in 2050	52.57	89.0	19
2050	Access in 2035	52.96	89.1	19
2050	Access in 2035 with import	44.23	81.3	0

5.9.1 The cost of a locked supply

Keeping the present supply costs 45.39 M€/yr more in 2035, and 70.85 M€/yr more in 2050. That is the distance between Diesel only and Access in 2050. The two serve the same households with the same demand. They differ only in whether the model may invest. The gap grows by 56 % between the two horizons.

The emissions go the same way. Investing avoids 210.1 ktCO₂eq per year in 2035 and 291.7 in 2050. Diesel only burns 689 GWh of fuel in 2035 and 945 in 2050, against 20 and 19 for the trajectory that invests.

Diesel only pays for fuel and for almost nothing else. On the electric scope it pays 77.8 M€/yr of resources in 2035, against 4.3 M€/yr of investment and maintenance. In 2050 it pays 107.0 against 5.4. Figure 26 gives the same picture for the seven runs.




Figure 26: Annual cost of the seven projection runs, split between investment, maintenance and resources. The hatched part of the investment is the capacity inherited from 2035 and charged again at its own cost, as described in Section 3.6.4. Grid extension is excluded.

5.9.2 The price of the calendar

Access in 2035 and Access in 2050 differ by 5.23 M€/yr at the 2035 horizon, and by 0.396 M€/yr at the 2050 one.

At 2035 the two do not serve the same people. Access in 2035 has connected everyone and Access in 2050 has not, so Access in 2050 also delivers less electricity, 278 GWh against 314. The 5.23 M€/yr buys a higher level of service at the same time as it costs more.

At 2050 the two are the same system. They hold the same households, the same split between the three sources and the same demand.

The two calendars also differ in what they deliver in between. The number of households with electricity is known in 2024, 2035 and 2050, and a linear interpolation between those three points gives a value for every year. Summed from 2025 to 2050, Access in 2035 delivers 2,648,871 household-years of access and Access in 2050 delivers 2,566,077. The difference is 82,793 household-years.

The lines are left out of these costs, and what that hides can be bounded. At 2035 Access in 2035 has built the whole network, 4.19 M€/yr. Access in 2050 has built only the lines that reach the households connected by then, and the spatial layer does not price that partial network. The lines therefore add between zero and 4.19 M€/yr to the cost of arriving early, so the 5.23 M€/yr is a lower bound of that cost. At 2050 both have built the same network.

5.9.3 What the model builds

At 2035 the two trajectories build the same kind of system and differ in its size. Access in 2035 installs 265 MW of utility solar and Access in 2050 installs 236 MW, producing 357 and 319 GWh. Lithium storage follows the solar, at 348 and 309 MWh. The old diesel fleet stays installed in both

and hardly runs, at 7.8 and 6.6 GWh against the 185 GWh it delivers in 2025. Diesel only builds nothing at all, and its 88 MW of engines produce 252 GWh.

At 2050 the two trajectories end on the same fleet. Both install about 330 MW of utility solar producing 434 GWh, both hold 412 MWh of storage, and both keep under 8 MW of diesel running 6.4 GWh. The two solar fleets differ by 1 MW on 330.

The 0.396 M€/yr between them comes from one thing. Access in 2035 built 265 MW of that solar fifteen years earlier, against 236 MW for Access in 2050. Those 29 MW were bought at the 2035 price and not at the 2050 one, which is 56.15 against 42.86 €/kW of annualised investment. The product is 0.38 M€/yr, so almost the whole gap. Figure 27 shows the two paths and where they meet.




Figure 27: Regional utility solar capacity at the three horizons. The two access calendars end on the same fleet in 2050. The import scenario starts from the 2035 fleet of Access in 2035 and builds nothing on top of it.

Diesel only in 2050 holds 55.6 MW of diesel, no solar and no storage. It produces 346 GWh, and the region burns 945 GWh of fuel. Its capacity is smaller than the 88 MW of 2035 and its output larger, so the engines run much closer to their annual ceiling.

5.9.4 The Brazilian import

The last run gives the region access to Brazilian electricity in 2050, at the wholesale price of Section 4.8.1. It imports 113.8 GWh, 76.3 into C3 and 37.5 into C5, and pays 6.75 M€/yr for it. The annual cost falls from 52.96 to 44.23 M€/yr, a saving of 8.74 M€/yr or 16.5 %. Utility diesel disappears, and the only fuel left in the region is burnt by household generators.

The import replaces storage more than solar. The lithium fleet falls from 412 to 144 MWh, so nearly two thirds of it goes. The solar fleet falls from 330 to 263 MW, and that fleet is the one built in 2035: the model keeps what it inherits and adds no new panel. Electricity that arrives at every hour of the year does what a battery does, and nobody has to build it.

The emission factor of that electricity is not fixed by anything. The greenhouse gas limit of the model does not bind, so the factor never enters the optimisation, and the emissions can be recomputed afterwards on the imported energy. Three values are used. The central one, 0.0385 ktCO₂eq/GWh, is the operational factor of the Brazilian system published for 2023. The upper one, 0.2065, is the generic mix of the national catalogue. The lower one, 0.002, is Jirau and Santo Antonio without reservoir emissions. None of the three is a full life cycle factor.

Table 33: Regional emissions of the import scenario in 2050, in ktCO₂eq/yr, under three emission factors and three volumes of Brazilian electricity. Access in 2035 without the import emits 89.1.

Imported (GWh/yr)	Emission factor (ktCO ₂ eq/GWh)		
	0.002	0.0385	0.2065
92.3	78.8	82.2	97.7
113.8	77.2	81.3	100.5
185.7	72.1	78.9	110.1

Table 33 recomputes the regional emissions under the three factors, and the sign of the result changes with the factor. Under the catalogue factor the import makes the region dirtier, and importing more makes it worse: 97.7 at the lowest volume and 110.1 at the highest, against 89.1 with no import at all. Under the central factor the ranking turns over, and importing more makes the region cleaner: 82.2 falls to 78.9. The cost of the import is settled and its carbon balance is not.

5.9.5 Sensitivity of the 2050 results

Eight variants were run on the 2050 configurations. Each one changes a single assumption. Table 34 gives what each does to the annual cost.

Table 34: Sensitivity of the 2050 results. Each line is the change in annual cost against the central run of the same scenario.

Scenario	Variant	Change (M€/yr)
Diesel only	diesel life of 20 years instead of 30	+0.24
Import	price of 0.03 instead of 0.0593 M€/GWh	−4.22
Import	price of 0.09 instead of 0.0593 M€/GWh	+4.16
Import	availability of 92.3 instead of 113.8 GWh/yr	+0.61
Import	availability of 185.7 instead of 113.8 GWh/yr	−1.14
Access in 2050	storage maintenance at 2.5 % of capital per year	+4.11
Access in 2035	solar water heater blocked	+0.16
Access in 2050	solar water heater blocked	+0.15

The life of the diesel fleet carries nothing. Cutting it from thirty to twenty years costs 0.24 M€/yr in the scenario where diesel does all the work. The capacity, the dispatch and the emissions stay the same, and only the annuity moves.

The import wins even at a higher price. At 0.09 M€/GWh it still costs 48.4 M€/yr against 52.96 without it. The price sweep is almost symmetric, -4.22 and $+4.16$ M€/yr, because the model takes the whole volume it is offered whatever the price. What decides here is the volume allowed, not the arbitrage between two prices. That sweep also moves the price of every external electricity, the national grid included, since the parameter carries no regional dimension. That side effect works against the import and not for it. Raising the volume matters less on the cost, since offering 185.7 GWh instead of 113.8 saves 1.14 M€/yr. It matters on the fleet. The model takes every kilowatt hour it is offered in the three variants, the solar capacity stays at 263 MW in all of them, and only the battery moves: 199 MWh at the lowest volume, 144 at the central one, and none at all at the highest. The import and the storage are two ways of buying the same service.

The maintenance of the storage is the one variant that moves a headline figure. The catalogue charges 0.66 €/kWh a year to maintain utility lithium storage, which is 0.16 % of its capital cost, where the baseline used for the cost ratios of this work puts that figure at 2.5 % [34]. The baseline rule adds 4.11 M€/yr to Access in 2050, or 7.8 % of its annual cost, and the optimiser pays almost all of it instead of avoiding it. Diesel only holds no storage and does not move, so the cost of a locked supply falls from 70.85 to 66.74 M€/yr, a reduction of six %. The bias is one-sided, since every scenario that builds storage comes out low while the current system and Diesel only do not, so rebuilding the supply looks cheaper than it is against keeping it.

Blocking the solar water heater costs 0.16 M€/yr in 2050 under Access in 2035 and 0.15 under Access in 2050, three tenths of a % in both. The same test on universal sufficiency in 2025 costs 3.30 M€/yr, or 10.6 %. The collector fleet of 2050 is 2 MW_{th} against 48 in universal sufficiency, and the electric resistance already produces 6.2 GWh of heat there against 3.2 for the collector. The switch has happened before the test is run. Section 6.5 reads that collapse and Chapter 8 says what causes it.

Figure 28 places the three trajectories on one pair of axes, starting from the current system of 2025.




Figure 28: Annual cost and emissions of the three trajectories, anchored on the current system of 2025. Only the vertical gaps at a fixed horizon are compared, since the slope between two horizons mixes several changes at once. The import point uses the central emission factor.

6 Discussion

Chapter 5 reports what the model finds. This chapter reads those findings from the point of view of someone who has to decide, and it states what the model does not settle. It follows the order of the decisions themselves: what to do first, where to do it, who pays for it, whether to buy electricity from Brazil, and what the cheapest system depends on. A last section says which of these results would hold in another region.

6.1 Rebuilding the supply comes first, and the order has a price

The three steps of Table 31 rank the levers without ambiguity. Rebuilding the supply saves 24.36 M€/yr and no household changes what it consumes. Giving the sufficiency bundle to the households that have nothing costs 4.92 M€/yr. Bringing the connected households down to the same bundle saves 15.79 M€/yr. The supply lever is the largest, it is the only one that requires nothing of the population, and it is the one a public operator can decide on its own.

Waiting makes it worse. Keeping the present supply costs 45.39 M€/yr more than rebuilding it in 2035, and 70.85 M€/yr more in 2050. The gap grows with the demand, because a locked supply pays for every additional kilowatt hour in fuel. Between the two horizons the resource cost of Diesel only rises from 77.8 to 107.0 M€/yr, while its investment and maintenance, the renewed engines included, rise from 4.3 to 5.4. Every year of delay is therefore paid in fuel, and the fuel bill grows with the region.

One of the three steps is easy to misread. Bringing the connected households down to the bundle measures a distance. It does not propose to cap anyone. What it shows is that the level of service assumed for the connected households drives the cost more than the access does.

The order itself also carries a cost. The same access costs 4.92 M€/yr once the supply has been rebuilt and 2.51 M€/yr on the diesel system of today. The reason is in Section 5.8: the old fleet is fully written off, it runs at night, and it firms the system at no capital cost. A planner should read that as an asset with an expiry date. The firming is free only as long as the engines are still turning and still amortised, and both stop being true during the period this thesis covers.

That does not reverse the ranking. The supply lever saves 24.36 M€/yr and the penalty it creates is 2.42 M€/yr, ten times less. What it changes is the budgeting. If the connections are planned after the diesel fleet has been switched off, part of their cost is firming that nobody had to pay for before, and it appears as batteries or as restarted engines rather than as lines. The cost of the access therefore cannot be budgeted on its own, since it depends on the state of the supply at the moment the connections are made. The model cannot arbitrate between the two orders, because it optimises each year on its own and never prices a sequence.

6.2 The cost and the access are not in the same municipalities

C3 and C5 carry between 69 and 82 % of the regional cost in every scenario, because Riberalta, Guayaramerín and Cobija are where the diesel is burnt. The households without a supply are somewhere else. Nearly half of the households of C4 have no public service, and C1 is in the same position. C5 is already 96 % connected.

The two programmes therefore do not compete for the same territory, and they are not the same kind of investment either. Rebuilding the supply is a few large sites in two clusters, and more than half of the saving is made in C3 alone. Reaching the unconnected households is 3,948 km of medium voltage line towards 626 scattered communities, and its cost is spread across the region. One is concentrated and one is dispersed, so they can run in parallel and the delay of one does not block the other. Another region with a different geography would not necessarily allow this.

The demography adds a rule that is easy to get wrong. The number of households of the region grows by 42.7 % by 2050 while the population grows by 20.0 %. Five municipalities lost inhabitants between 2012 and 2024 and gained households anyway, and they hold a fifth of the households of the region. Connections follow households, so a municipality that is losing people is not a municipality that can be deprioritised. Any programme sized on population projections will undercount what it has to build.

6.3 The thesis prices the system, not the bill

Every cost in this thesis is a system cost. It is what the region consumes in fuel, capital and maintenance, and it is not what anyone actually pays. Two gaps separate the two. The model buys diesel at its border price of 0.1113 M€/GWh, while the pump price in Bolivia is about half of that, and Section 4.8.1 explains that the cost of hauling fuel to the Amazon partly closes the difference without measuring by how much. The supply cost of the current system is 0.33 €/kWh.

The first gap can be bounded with the numbers of Chapter 5. The current system burns 512 GWh of diesel a year and the supply transition burns 34. Valued at the pump price, the fuel bill of the first falls by about 29 M€/yr and the fuel bill of the second by about 2. The 24.36 M€/yr that the supply transition saves therefore disappear before the supply is even re-optimised for the cheaper fuel. The

saving is real for the country, which imports the diesel and pays the difference between the two prices, and it is invisible to the utility, which buys at the pump. Those 29 M€/yr are an upper bound of what the subsidy costs the state in this region alone, since part of the gap between the two prices is the cost of hauling the fuel and not a subsidy. Even halved, that figure is more than twice the cost of universal access. Both gaps say the same thing: part of what the present system costs is paid outside the region.

That matters for the transition, because it changes shape rather than size. The current system pays 62.80 M€/yr of resources and almost no capital. The supply transition pays 28.49 M€/yr of annualised investment and almost no fuel. The second is cheaper in total, but it has to be financed before the first euro of fuel is saved, and a recurring fuel bill and a capital requirement are not funded by the same institutions.

The orders of magnitude are not out of reach. The IDTR III programme commits 148 million USD to about 25,000 grid connections and 9,000 solar home systems nationally, and closes in 2030 [61]. The 37.63 M€ of medium voltage line that this region needs is a fraction of that. At national level about 60,000 households are connected each year [81], so the 20,714 households of Sources B and C are about four months of the national effort, spread over a region that holds under three % of the country. Nothing here says the money is unavailable. What it says is that the constraint is not the total.

This thesis does not say who pays, at what tariff, or through which instrument. That question is the real difference between the present system and the transition, and it is outside the model rather than inside it.

6.4 The Brazilian import is cheap, and everything else about it is undecided

Buying 113.8 GWh from Brazil in 2050 removes 8.74 M€/yr from the annual cost, which is 16.5 %. It survives its own price: at 0.09 M€/GWh, half again above the wholesale price used, the import still wins. On cost alone the case is not close.

What the import replaces matters more than what it saves. The solar fleet falls by a fifth, but the lithium fleet falls from 412 to 144 MWh. Electricity available at every hour of the year does what a battery does, and it does it without being built. The interconnection decision and the storage decision are therefore the same decision, taken years apart. The region cannot size its storage correctly without knowing whether the line will arrive, and 268 MWh of lithium is what the answer is worth in this scenario. Offering more makes the point sharper still: at 185.7 GWh the model builds no battery at all.

That exposure is visible here because the model has no foresight. It invests in 2035 as if the line would never come, and the interconnection then arrives in 2050 into a system built without it. The whole solar fleet of 2035 stays standing and none of it is added to. Placing the import in 2050 makes this the mildest version of the problem, since that fleet has fifteen years of undisturbed operation first. An interconnection that arrived between the two horizons would strand the same fleet with fewer years of use behind it. That case is not run here.

What the line strands is the storage more than the panels. In the import scenario the model keeps 263 of the 330 MW of utility solar in service and retires two thirds of the batteries. Whether it would build that solar from scratch with the import already available is not tested here, since the panels it keeps are inherited from 2035. Two facts still hold. Solar is cheaper than the diesel it replaces today, at the border price, and the national grid has been announced for this region since 2012 without a date. Waiting for a line before building panels is what the region has done for more than a decade, and the diesel bill is what it has cost. The decision that should wait for the line is the size of the storage, not

the panels.

The carbon result does not follow the cost result. The region emits 81.3 ktCO₂eq per year under the operational factor of the Brazilian system and 100.5 under the generic catalogue factor, against 89.1 with no import at all. The import can therefore make the region cleaner or dirtier, and which one it does depends on where the electricity comes from. Nothing secures that origin. The 2024 memorandum is a framework with no volume, no price, no date and no duration [36], and the loan proposal describes a purchase made on the Brazilian open contracting market, where public and private suppliers trade alike [37]. Neither document names a plant, so the hydropower usually cited in the public discussion is not reserved to anyone.

Two consequences follow for a planner. An import contract that specifies an origin is worth something that this thesis can price, and one that does not is worth nothing on the carbon axis. A region that covers a large share of its supply from a neighbour with no contract has taken on a dependency that a cost model does not see. This model resolves cost and energy balances, not security of supply, so nothing here says what happens to Riberalta or Cobija on a day when the line is unavailable.

The domestic alternative to the interconnection is also still open. The extension of the national grid towards Pando has been in pre-investment since 2012 with no financing secured and no construction started [35]. Neither route has a date, and the region has been planning around both for more than a decade.

6.5 The cheapest system needs a solar collector on every roof

The universal sufficiency result rests on one technology more than on any other. Blocking the solar water heater raises the annual cost by 3.30 M€/yr, which is four fifths of the 3.99 M€/yr that the entire medium voltage network costs.

The technology itself is not in doubt in this climate. A thermosiphon delivered hot water on more than nine days out of ten in a Manaus field trial [65], a Bolivian supply chain existed under the El Sol project [66], and Brazilian social housing has been equipping roofs with collectors at scale for years [67]. What none of that establishes is a deployment of the size this scenario implies, which is a collector on essentially every roof of 21 municipalities. The cost of that deployment is in the model. The installers, the spare parts and the maintenance network are not.

This lever is easier to act on than any other in this thesis, because it does not depend on the grid. A dispersed household with a standalone system needs the same collector as a connected one, which is why the collector cancels out of the breakeven of Section 4.9.2. Every other lever waits for a line or for a power plant. This one can start now in any municipality, and it reduces the electricity the system has to build later.

The result should not be read past 2025. Section 5.9 shows the same test costing 0.16 M€/yr in 2050. Two things cause that collapse. The projections carry a hot water demand eight times smaller, so the collector fleet is 2 MW_{th} instead of 48. Photovoltaic capital also falls to 46 percent of its value by 2050 while the collector keeps its 2025 price for want of a source. The 2050 figure says nothing about solar thermal collectors and Chapter 8 reports it as such.

6.6 What would hold in another region

The methodological contribution is the most portable part of this work. The absolute bounds on annual production of Section 3.4 are needed wherever one population must stay exactly as it is while

another is served, which is the normal shape of an access question. Nothing in them is specific to Bolivia or to this model.

Two results depend on conditions that are common but not universal. Access costing more on a decarbonised system than on the system it replaces requires a thermal fleet that is written off and still running, which describes most isolated regions of the Global South today and fewer of them as those fleets retire. The separation of the population into three groups requires a census that records the source of electricity rather than its presence, which many do and Bolivia does since 2012.

One result does not transpose. The sufficiency ceiling is worth more than the access floor here because the region is very unequal inside itself. In a region where every household consumes about the same, capping the demand of the connected ones releases nothing, and the ranking of the three steps would change. What transfers is the decomposition itself, and each region has to compute its own numbers.

The resolution is a result of its own. A national model represents the rural lowlands of Bolivia as one block and concludes that standalone systems are the near-term solution for dispersed households [3]. At municipal resolution the answer splits: 626 of the 697 communities of this region are cheaper to reach with a line today and 71 are not, and the dispersed ones sit almost entirely in C1 and C4. Whether a community is better served by a line or by a solar home system depends on that community, and it can only be seen at a resolution where communities exist.

7 Conclusion

This thesis measures the cost and the emissions of alternative energy transition pathways towards universal access to electricity in the Bolivian Northern Amazon. Most of what a planner needs did not exist and had to be built first. RAMP, a stochastic model, produced the demand of the households that no meter records. EnergyScope MultiCell, a linear optimisation of the whole energy system, sized the supply that serves it. A spatial layer classified the 697 communities of the region between a line extension and a standalone system. Five scenarios cover 2025, and seven further runs cover 2035 and 2050.

Three steps separate the system the region operates today, at 66.51 M€/yr, from one that gives every household the sufficiency bundle, at 31.28. Replacing diesel generation with utility solar and storage saves 24.36 M€/yr, 37 percent of the annual cost, and no household changes what it consumes. Serving the 20,714 households that are badly served or not served at all then costs 4.92 M€/yr. Bringing the households that are already connected down to the same bundle saves another 15.79 M€/yr, because those households consume more than the bundle today. The supply is therefore the largest lever, and the only one that asks nothing of the population. The level of service assumed for the connected households is the second largest, three times heavier than access. The comparison measures a distance between two assumptions and recommends capping no one.

The generation optimum does not see all of that cost. EnergyScope assumes that every household can already be reached, so the price of reaching it is invisible to the model. The spatial layer supplies it: 626 of the 697 communities are cheaper to reach with a medium voltage line than to equip with a standalone system, and those lines cost 3.99 M€/yr, annualised over thirty years. Universal access therefore costs 6.50 M€/yr more than the current system, and only 2.51 of that is generation. The price of access also depends on the supply it is added to: on a supply already rebuilt with solar, the same access costs 4.92 M€/yr of generation instead of 2.51, because the old engines are already paid for and cover the night today.

The emissions answer is simpler. The two scenarios that keep the diesel plants emit around 217

ktCO₂eq per year, and the three that replace them emit between 58 and 74, although the demand they serve differs widely. Giving access without touching the plants leaves emissions at 217.3 against 217.7, because the new demand is served by new solar. Emissions depend on whether the plants are replaced, and on almost nothing else. The projections then price two decisions. The first is whether the region invests at all. By 2035 the annual cost of a region that keeps its diesel plants exceeds that of one that rebuilds its supply by 45.39 M€/yr, and by 2050 by 70.85, because the demand grows. The second is the year in which every household is reached. Two trajectories were run, one reaching universal access in 2035 and the other in 2050. Both end on the same fleet, about 330 MW of utility solar, and arriving early costs 5.23 M€/yr more in 2035 and 0.396 more in 2050. Cost does not decide the pace. The rate at which lines can be built does, about 372 km a year to reach every connectable community by 2035, half of what ENDE reports building across four departments in its best documented year.

Three recommendations follow. Rebuilding the supply comes first: it is the largest saving, the largest emission cut, and the only lever a public operator can pull without changing anything in a household. The access programme should be planned together with the supply and not after it, because the two reach different municipalities and do not compete for the same sites, and because connecting households while the amortised engines still cover the night costs less than buying that night-time supply again later. Finally, no decision on storage should be taken without a position on the Brazilian interconnection: an import available at every hour covers the same hours a battery would, and buying 113.8 GWh in 2050 removes 268 MWh of battery storage and 8.74 M€/yr from the optimal system.

Two things in this work are reusable beyond the region. EnergyScope MultiCell could only bound a technology as a share of a quantity the optimiser itself decides. That fails as soon as one population must stay unchanged while another is served. The floor and the ceiling on annual production added here remove that limit. And the comparison itself has no precedent on one territory: the sufficiency literature carries no measured counterfactual, and building both here makes the three steps add up exactly. These numbers still have limits. They are system costs. They say what the region consumes in fuel, capital and maintenance, and not who pays for it: at the subsidised pump price the saving of the supply transition accrues to the state and not to the utility. The model balances energy and cost, not frequency or reserve. The cheapest system also assumes a solar collector on almost every roof, proven in this climate but never deployed here at that scale. What the region can do is nonetheless clear: it can serve every one of its households, including the 20,714 that are served badly or not at all, with a system that costs less than the one it operates today and emits a quarter as much.

8 Limitations and future work

The limitations of this work are stated where they arise, next to the choice or the input that carries them. This chapter keeps the ones that weigh most on the results, and says for each what it does to the numbers.

Some limits come from the model itself. Every consumer of a cluster draws from one shared electricity layer. The production bounds added in this work fix how much energy a home system delivers over the year, but not which household receives it, so every statement about home systems concerns volumes and not topology. The model also closes an energy balance and nothing else. The year is reduced to 16 typical days and storage returns to its starting level at the end of each day, so a week without sun does not exist in the model, and nothing here concerns frequency or reserve. Both simplifications work in the same direction: the systems that rebuild the supply come out smaller and cheaper than a system built to a reliability standard would. Domestic hot water has no hourly profile, which favours the solar collector against the electric resistance. Finally, the spatial layer knows two

options, a line or a household system. The village mini-grid, the option the Bolivian literature has studied most, is absent, and part of the 71 dispersed communities would be served differently with it.

Other limits come from the data. The duty cycle of the refrigerator is the assumption the results are most sensitive to, and almost no data exists behind it: those refrigerators draw two thirds of everything delivered to an end use in universal sufficiency. The connected households have no measured load curve, so the shape of 98 percent of the regional demand is borrowed from the 2 percent that RAMP models directly. The technology catalogue also charges 0.16 percent of the capital cost per year to maintain utility storage, where the NREL baseline used for the cost projections of this work puts that figure at 2.5 percent. The catalogue value was kept, because every technology of the model comes from the same database and rebasing one parameter alone would have meant rerunning every scenario. The gap is worth 4.11 M€/yr on one scenario, and it makes every system that builds storage look cheaper than it is.

Two further biases have a known direction. The network technologies carry no construction emissions in the database. The solar panels and the batteries do carry theirs, but the 3,948 km of medium voltage line, the inter-cluster lines and the hot water tanks carry none, so every scenario that connects households emits more than reported here. The solar collector keeps its 2025 price while photovoltaic capital falls by half, because no source gives its decline. This is why the water heating lever collapses between 2025 and 2050, and that collapse says nothing about the technology itself.

8.1 Future work

The work carried out in this thesis could be extended in several ways.

- The two weakest assumptions of the demand side could both be replaced by measurement. A campaign on a few dozen refrigerators would give the duty cycle that sizes the whole system, and a recorded load curve on one isolated system would give the shape that is currently borrowed from the off-grid households.
- The dispersion constraint could be repaired. An alternative form already sits in the model code, commented out, with the exchange terms in its denominator, and testing it removes the weakness that leaves the constraint inert in any importing cluster.
- The Brazilian import is modelled at a single date and without foresight, so the region invests in 2035 as if the line would never come. Running the import at several dates, with the 2035 decision anticipating it, would separate the cost of the import from the cost of being surprised by it, and the 268 MWh of storage it makes useless is a lower bound on that effect.
- A systematic sensitivity analysis could replace the handful of levers tested here. The discount rate, the diesel price and the rating of the household generators deserve it most, and a robust formulation, which optimises against a range of futures rather than a single one, is the natural step after that.
- Full sector coupling, adding transport and industry to the power to heat coupling already present, would let the region be analysed as an energy system and not as an electricity system with a heat layer attached.

Appendices

A Clustering details

Table 35: Variables used for municipality clustering.

Variable	Description	Source
Population density	Inhabitants per km ²	Census 2024 [1]
GDP per capita	USD/person	Andersen et al. [5]
Cp_PV	PV capacity factor	Renewables Ninja [42, 82]
Cp_Wind	Wind capacity factor	Renewables Ninja [42, 83]
Temperature	Annual mean (°C)	Renewables Ninja [42]
Altitude	Metres above sea level	SRTM DEM [84]
Relative humidity	Annual mean (%)	ERA5 reanalysis [85]
Num. vehicles	(Cars + motorcycles) / population	Census 2024 [1]
EL_network	Binary: SIN physically present	Geoportal [38]
FG_network	Binary: gas pipeline present	ANH [39]
EL_demand	Measured consumption (MWh/pers/yr)	Andersen et al. [5]





Figure 29: Normalised weights of the twelve clustering variables, from the seven survey responses.

Table 36: Survey results: importance scores assigned by seven participants (P1–P7) and resulting normalised weights.

Variable	P1	P2	P3	P4	P5	P6	P7	Weight
EL_network_access	15	13	30	10	20	25	12	0.179
FG_network_access	10	9	10	5	20	25	10	0.127
EL_demand	10	10	5	7	10	6	1	0.070
Population density	10	13	10	15	10	10	18	0.123
GDP per capita	5	9	20	5	10	5	12	0.094
Vehicles per capita	5	9	5	5	5	5	6	0.057
Cp_Wind	10	7	0	5	5	5	6	0.054
Cp_PV	10	5	0	10	5	7.5	2	0.056
Altitude	3	4	5	10	2	1.5	5	0.044
Temperature	4	2	5	5	2	2.5	1	0.031
Relative humidity	3	3	5	10	1	1.5	4	0.039
Geographic distance (α)	15	16	5	13	10	6	23	0.126

Figure 30: Cluster selection metrics for $k = 2$ to $k = 21$.

B Demand inputs and RAMP profiles

The sufficiency demand was generated for all 21 municipalities. Two of them are shown here to illustrate the outputs, Santos Mercado with 601 households and Ixiamas with 3,306.

Figure 31 gives the full year demand of Santos Mercado, stacked by sector. Households dominate all year. The contribution of the income generating activities, which includes rice processing, is visible in summer and fall and disappears in winter and spring. Community services form a thin constant layer. The slight rise in household demand during winter and spring comes from the longer fan usage times of those seasons, calibrated in Section 4.5.1.

Table 37: Infrastructure allocation rules for the BNA (sufficiency demand). n is total_persons and hh is total_hh.

Infrastructure	Rule	Total (21 mun.)
School (big_school)	$\max(1, n//5,000)$	62
Health centre	$\max(1, n//5,000)$	62
Public lighting	$hh//10$	8,412
Store	$\max(1, hh//30)$	2,796
Restaurant	$\max(1, hh//30)$	2,796
Workshop	$\max(1, hh//60)$	1,393
Entert. business	$\max(1, hh//60)$	1,393
Rice processing	$\max(1, hh//200)$	413

Table 38: Household-level appliance inputs for the sufficiency demand.

End-use	Appliances	Seasonal behaviour
Lighting	LED interior + exterior	Shorter use in summer, longer in winter
ICT	1 TV, 1 radio (3 W), 1 laptop, 4 phone chargers	No seasonal variation
Cold storage	1 refrigerator	Higher duty cycle in warm seasons
Thermal comfort	1 fan per household	All 4 seasons; per-municipality calibration (Sec. 4.5.1)
Water heating	Electric, groundwater temp. model	Higher demand in colder months




Figure 31: Full year community demand for Santos Mercado, stacked by sector.

Table 39: Rice processing equipment specifications.

Equipment	Power (W)	Daily use (min)
Rice dryer	2,500	480
Rice huller	2,000	300
Rice polisher	1,500	300
Sorter/grader	1,000	300
Packaging	400	180

Table 40: Mapping of RAMP output categories to EnergyScope demand parameters.

RAMP category	EnergyScope parameter	Sector
Illumination (all)	LIGHTING_R_C	HOUSEHOLDS / SERVICES
Public lighting	LIGHTING_P	PUBLIC_LIGHTING
ICT, water supply, medical equip.	ELECTRICITY	HOUSEHOLDS / SERVICES
Cold storage (all)	FOOD_PRESERVATION	HOUSEHOLDS / SERVICES
Thermal comfort (fans, AC)	SPACE_COOLING	HOUSEHOLDS / SERVICES
Water heating (households, health centres)	HEAT_LOW_T_HW	HOUSEHOLDS / SERVICES
Kitchen (restaurants)	COOKING	SERVICES
Workshop machinery	MECH._ENERGY_COMM	SERVICES
Rice processing	MECH._ENERGY_IND	INDUSTRY

Figure 32 compares the annual load duration curves of the two municipalities. Ixiamas peaks at about 250 kW and Santos Mercado at about 57 kW. Both curves have the shape expected of a residential demand, with a moderate ratio between the peak and the base and no spike that would indicate a large industrial load.




Figure 32: Annual load duration curves for Ixiamas and Santos Mercado, sufficiency demand.

C EnergyScope configuration tables

Table 41: Inter-cluster distances (km), centroid-to-centroid (Haversine).

	C1	C2	C3	C4	C5
C1	–	384.4	328.9	304.3	385.0
C2		–	283.3	113.1	15.4
C3			–	179.2	270.7
C4				–	104.4
C5					–

Although mobility demand is zero in the current scenario, the modal shares are kept for consistency with the national model and to ease future extensions that include transport.

Table 42: Key resource availability parameters by cluster.

Resource	Allocation	Source and notes
WOOD (local only)	Proportional to area	Base: SA-PA = 58.39 GWh/yr [86] via [30]. avail_exterior = 0 in all clusters; wood is not imported
BIOMASS_RESIDUES (local)	Proportional to area	Base: SA-PA = 226.76 GWh/yr [86]
DIESEL (external)	Per cluster	Opportunity price: 0.1113 M€/GWh (border price [29]). Delivered at that price in every cluster (Section 4.8.2).
LPG (external)	Per cluster	Subsidised price 0.0253 M€/GWh (YPFB garrafa 22.50 Bs/10 kg, FX 6.96 Bs/USD)
ELECTRICITY (external)	C1 only, non-binding	Set above the Source A grid consumption of the C1 municipalities, 11.593 GWh/yr (Reyes, Santa Rosa Beni, Ixiamas at 100% SIN; Exaltación at 50% SIN). Price: 0.0593 M€/GWh (CNDC precio monómico). All other clusters: 0
GAS (local)	0 everywhere	No local gas fields in the BNA
LNG (external)	1 GWh/yr per cluster	The only path by which natural gas enters the system, through a regasification step onto the GAS layer. Price 0.1988 M€/GWh, inherited unchanged from the national catalogue

D Projection details




Figure 33: Deviation of the predicted 2024 cluster population from the census, for the six candidate

Table 43: Key per-cluster configuration parameters.

Parameter	C1	C2	C3	C4	C5	Notes
import_capacity	2	2	2	2	2	Big M; real constraint is freight cost
solar_area_*	proportional to area					Base SA-PA = 47,764 km ² [30]
solar_area_high_irr	0	0	0	0	0	No high-irradiance zones
share_heat_dhn	0	0	0	0	0	No district heating
share_freight_train	0	0	0	0	0	No railways
share_short_haul_flights	0	0	3%	0	3%	Airports in Riberalta, Guayaquerín, Cobija
share_mobility_public	about 72.9%					Energy balance 2021 [87], disaggregated by population
share_freight_boat	about 2.1%					idem

The network to build is concentrated in two clusters, and it barely grows between the three horizons.




Figure 34: Medium voltage line to build in each cluster, at the household counts of the three horizons.

Table 44: Treatment of each appliance in the projection of Source A ownership.

Appliance	Treatment	Reason
Refrigerator, air conditioning, microwave, washing machine	access relationship	Explained variance above 0.3, and freezing is physically impossible: San Pedro would hold no refrigerator in 2050 at full access
Computer, internet, telephone	access relationship	Favoured by the out of sample test
Radio, television	frozen	The test favours freezing, and ownership is already close to saturation
Water heater	frozen	Explained variance of 0.29; access does not explain ownership
Antenna	follows the radio	Already shares its penetration
Fan	unchanged	Comes from the thermal model of Section 4.5.1, not from the census

E Additional figures




Figure 35: Households of the BNA by source of electricity, Census 2024.

Table 45: Households by source and by cluster, given as Source A / Source B / Source C. The two trajectories share the same 2024 baseline and converge on the same 2050 state. Column totals may differ by one household from Table 19, since each column is rounded on its own.

Cluster	2035			
	2024	Access in 2035	Access in 2050	2050, both
C1	5,843 / 2,656 / 2,434	12,819 / 124 / 0	9,862 / 1,660 / 1,420	14,727 / 146 / 0
C2	156 / 393 / 253	998 / 0 / 0	655 / 208 / 134	1,204 / 0 / 0
C3	33,926 / 2,274 / 4,098	47,877 / 81 / 0	44,082 / 1,398 / 2,478	55,551 / 90 / 0
C4	8,649 / 3,754 / 4,209	20,574 / 227 / 0	16,506 / 2,095 / 2,200	25,280 / 207 / 0
C5	14,921 / 248 / 395	19,084 / 0 / 0	18,861 / 86 / 137	22,994 / 0 / 0




Figure 36: The extension capital, split between the three posts of Equation (15).

References

- [1] Instituto Nacional de Estadística (INE), *Censo de población y vivienda 2024 – tabulados por municipio y TIOC*, Accessed: 2026, 2025. [Online]. Available: <https://cpv2024.ine.gob.bo/index.php/tabulados-sobre-la-tematica-pobreza/>.
- [2] ABC – Administradora Boliviana de Carreteras, *Plan lluvias 2024–2025*, Plan addressing climate-related road disruptions on the Red Vial Fundamental. Accessed: 2026, 2024. [Online]. Available: <https://ahoraelpueblo.bo/index.php/nacional/economia/la-abc-inicia-el-plan-lluvias-para-garantizar-rutas-transitables-en-la-red->.
- [3] R. Arias Valentines, “Integrating isolated communities into national energy planning models – a case study in Bolivia”, M.S. thesis, Université de Liège, 2025. [Online]. Available: <https://matheo.uliege.be/handle/2268.2/23806>.
- [4] IEA, IRENA, UNSD, World Bank, and WHO, “Tracking SDG 7: The energy progress report 2025”, World Bank, Washington, DC, Tech. Rep., 2025. [Online]. Available: <https://tracking7.esmap.org/>.
- [5] L. E. Andersen, B. Branisa, and F. Calderón, “Estimaciones del PIB per cápita y de la actividad económica a nivel municipal en Bolivia en base a datos de consumo de electricidad”, CIS, La Paz, Tech. Rep., 2019.
- [6] M. Bhatia and N. Angelou, “Beyond connections: Energy access redefined”, Energy Sector Management Assistance Program (ESMAP), World Bank, Washington, DC, Tech. Rep. 008/15, 2015. [Online]. Available: https://mtfenergyaccess.esmap.org/data/files/download-documents/full_report_beyond_connection.pdf.
- [7] P. J. Zabalaga et al., “Towards a sustainable Bolivian energy system in 2050: The pathway for decarbonization under high renewable potential”, *Energy*, vol. 337, p. 138 335, 2025. DOI: 10.1016/j.energy.2025.138335.
- [8] N. Stevanato, L. Rinaldi, S. Pistolese, S. L. B. Subieta, S. Quoilin, and E. Colombo, “Modeling of a village-scale multi-energy system for the integrated supply of electric and thermal energy”, *Applied Sciences*, vol. 10, no. 21, p. 7445, 2020. DOI: 10.3390/app10217445.
- [9] S. Balderrama, F. Lombardi, F. Riva, W. Canedo, E. Colombo, and S. Quoilin, “A two-stage linear programming optimization framework for isolated hybrid microgrids in a rural context: The case study of the “El Espino” community”, *Energy*, vol. 188, p. 116 073, 2019. DOI: 10.1016/j.energy.2019.116073.
- [10] S. Mandelli, M. Merlo, and E. Colombo, “Novel procedure to formulate load profiles for off-grid rural areas”, *Energy for Sustainable Development*, vol. 31, pp. 130–142, 2016. DOI: 10.1016/j.esd.2016.01.005.
- [11] F. Lombardi, S. Balderrama, S. Quoilin, and E. Colombo, “Generating high-resolution multi-energy load profiles for remote areas with an open-source stochastic model”, *Energy*, vol. 177, pp. 433–444, 2019. DOI: 10.1016/j.energy.2019.04.096.
- [12] F. Lombardi et al., “RAMP: Stochastic simulation of user-driven energy demand time series”, *Journal of Open Source Software*, vol. 9, no. 98, p. 6418, 2024. DOI: 10.21105/joss.06418.
- [13] C. L. S. Solis, P. D. Betta, N. Stevanato, S. L. B. Subieta, E. Colombo, and S. Quoilin, “Electricity demand forecasting for rural communities in developing countries: Calibrating a stochastic model for the Bolivian case”, in *Proceedings of ECOS 2023 – The 36th International Conference on Efficiency, Cost, Optimization, Simulation and Environmental Impact of Energy Systems*, Las Palmas de Gran Canaria, Spain, 2023. DOI: 10.52202/069564-0176. [Online]. Available: <https://www.proceedings.com/content/069/069564-0176open.pdf>.

- [14] C. Sanchez-Solis, S. Balderrama, G. Crevani, N. Stevanato, and S. Quoilin, “Bridging the energy gap: A Global South perspective on achieving universal access through energy sufficiency”, *Energy for Sustainable Development*, vol. 88, p. 101 793, 2025. DOI: 10.1016/j.esd.2025.101793.
- [15] T. Moss et al., “The modern energy minimum: The case for a new global electricity consumption threshold”, Energy for Growth Hub, Tech. Rep., 2020.
- [16] F. Riva, F. Gardumi, A. Tognollo, and E. Colombo, “Soft-linking energy demand and optimisation models for local long-term electricity planning: An application to rural India”, *Energy*, vol. 166, pp. 32–46, 2019. DOI: 10.1016/j.energy.2018.10.067.
- [17] S. Balderrama, F. Lombardi, N. Stevanato, G. Peña, E. Colombo, and S. Quoilin, “Surrogate models for rural energy planning: Application to Bolivian lowlands isolated communities”, *Energy*, vol. 232, 2021. DOI: 10.1016/j.energy.2021.121108.
- [18] S. L. Balderrama Subieta, “Optimal design and deployment of isolated energy systems: The Bolivian pathway to 100% rural electrification”, Ph.D. dissertation, Université de Liège, 2022. [Online]. Available: <https://hdl.handle.net/2268/267954>.
- [19] S. Pfenninger, A. Hawkes, and J. Keirstead, “Energy systems modeling for twenty-first century energy challenges”, *Renewable and Sustainable Energy Reviews*, vol. 33, pp. 74–86, 2014. DOI: 10.1016/j.rser.2014.02.003.
- [20] M. Howells et al., “OSeMOSYS: The open source energy modeling system. an introduction to its ethos, structure and development”, *Energy Policy*, vol. 39, no. 10, pp. 5850–5870, 2011. DOI: 10.1016/j.enpol.2011.06.033.
- [21] G. Limpens, S. Moret, H. Jeanmart, and F. Maréchal, “EnergyScope TD: A novel open-source model for regional energy systems”, *Applied Energy*, vol. 255, p. 113 729, 2019. DOI: 10.1016/j.apenergy.2019.113729.
- [22] P. Thiran, “Exploring options for a fossil-free European energy system: The role of renewable fuels”, Ph.D. dissertation, UCLouvain, 2024. [Online]. Available: <https://dial.uclouvain.be/pr/boreal/object/boreal:293132>.
- [23] C. A. A. Fernandez Vazquez, P. Jimenez Zabalaga, S. Balderrama, E. Cardozo, S. Quoilin, and H. Jeanmart, “A bi-directional soft-linking method for a whole energy system model and a power system optimization model: Application and analysis for the Bolivian case”, *Environmental Research: Energy*, vol. 3, no. 2, p. 025 011, 2026. DOI: 10.1088/2753-3751/ae5aa2.
- [24] D. Mentis et al., “Lighting the World: The first application of an open source, spatial electrification tool (OnSSET) on Sub-Saharan Africa”, *Environmental Research Letters*, vol. 12, no. 8, p. 085 003, 2017. DOI: 10.1088/1748-9326/aa7b29.
- [25] J. G. P. Balderrama et al., “Incorporating high-resolution demand and techno-economic optimization to evaluate micro-grids into the open source spatial electrification tool (OnSSET)”, *Energy for Sustainable Development*, vol. 56, pp. 98–118, 2020. DOI: 10.1016/j.esd.2020.02.009.
- [26] N. Stevanato, S. Corigliano, M. Petrelli, F. Tonini, M. Merlo, and E. Colombo, “Rural areas electrification strategies through shadow costs analysis: Bolivian Highlands case study”, *Energy for Sustainable Development*, vol. 65, pp. 162–174, 2021. DOI: 10.1016/j.esd.2021.10.009.
- [27] N. Moksnes, A. Korkovelos, D. Mentis, and M. Howells, “Electrification pathways for Kenya: Linking spatial electrification analysis and medium to long term energy planning”, *Environmental Research Letters*, vol. 12, no. 9, p. 095 008, 2017. DOI: 10.1088/1748-9326/aa7e18.

- [28] I. Pappis et al., “Influence of electrification pathways in the electricity sector of Ethiopia: Policy implications linking spatial electrification analysis and medium to long-term energy planning”, *Energies*, vol. 14, no. 4, p. 1209, 2021. DOI: 10.3390/en14041209.
- [29] CIE-UMSS, *EnergyScope MultiCell Bolivia – GitHub repository*, Accessed: 2026, 2024. [Online]. Available: https://github.com/CIE-UMSS/EnergyScope_Multicell_B0.
- [30] P. J. J. Zabalaga, X. Rixhon, J. Villarroel-Schneider, E. Cardozo, and H. Jeanmart, *EnergyScope Multicell Pathway: A multi-regional framework for long-term energy transition planning with application to Bolivia’s decarbonization*, Manuscript. Associated code and data: https://github.com/CIE-UMSS/EnergyScope_Multicell_Pathway_B0, 2025.
- [31] Instituto Nacional de Estadística (INE), *Proyecciones de población por municipio y TIOC, 2024–2034*, Official municipal horizon stops at 2034. Accessed: 2026, 2026. [Online]. Available: <https://nimbus.ine.gob.bo/index.php/s/Gx4TRNsi4fT2TH8/download>.
- [32] Instituto Nacional de Estadística (INE), *Población estimada y proyectada con tasas de crecimiento, 2000–2050*, National deceleration profile used to extend the regional growth rate beyond 2034. Accessed: 2026, 2026. [Online]. Available: <https://nimbus.ine.gob.bo/index.php/s/kdz7ktspsFz6C4/download>.
- [33] Instituto Nacional de Estadística (INE), “Proyecciones y estimaciones de población, revisión 2025 – documento metodológico”, INE, La Paz, Bolivia, Tech. Rep., 2026, Section 7 describes the municipal and TIOC method, including the backtest protocol on the 2001 and 2012 censuses. Accessed: 2026. [Online]. Available: <https://nimbus.ine.gob.bo/index.php/s/PaHRM4KgJXASM6n/download>.
- [34] National Renewable Energy Laboratory (NREL), *2024 annual technology baseline*, Moderate technology innovation scenario. Base year 2022. Accessed: 2026, Golden, CO, 2024. [Online]. Available: <https://atb.nrel.gov/electricity/2024/index>.
- [35] ENDE – Empresa Nacional de Electricidad, *Proyecto interconexión de Pando al SIN*, 1,025 km, 230 kV, Trinidad–Cobija. Pre-investment phase since 2012. Accessed: 2026, 2012. [Online]. Available: <https://www.ende.bo/proyectos>.
- [36] Agencia Boliviana de Información (ABI), *Hidroeléctricas brasileñas proveerán electricidad a municipios de Pando*, Santa Cruz, 9 July 2024. Accessed: 2026, Jul. 2024. [Online]. Available: <https://www.abi.bo/index.php/component/content/article/36-notas/noticias/economia/52493>.
- [37] Banco Interamericano de Desarrollo (BID), “Programa de inversiones para la transición energética del Norte Amazónico Boliviano – propuesta de préstamo”, BID, Tech. Rep. BO-L1238, 2026, Section 1.15 describes purchases on the Brazilian *Ambiente de Contratación Libre*; Section 1.26 states that public and private companies trading on that market may supply the Bolivian demand. Published 7 July 2026. Accessed: 2026. [Online]. Available: <https://www.iadb.org/document.cfm?id=EZIDB0001855-143724827-343>.
- [38] Ministerio de Hidrocarburos y Energías, Bolivia, *Geoportal – mapa del sistema interconectado nacional*, Accessed: 2026. [Online]. Available: <https://geoportal.mhe.gob.bo/maps/774/view>.
- [39] Agencia Nacional de Hidrocarburos (ANH), Bolivia, *Mapa de gasoductos troncales de Bolivia*, Accessed: 2026. [Online]. Available: https://www.anh.gob.bo/InsideFiles/Documentos/Documentos_Id-125-140319-27.pdf.
- [40] E. Schubert and P. J. Rousseeuw, “Fast and eager k-medoids clustering: O(k) runtime improvement”, *Information Systems*, vol. 101, p. 101 804, 2021. DOI: 10.1016/j.is.2021.101804.

-
- [41] CIE-UMSS, *RAMP Bolivia – GitHub repository*, Accessed: 2026, 2024. [Online]. Available: https://github.com/CIE-UMSS/RAMP_Bolivia.
- [42] S. Pfenninger and I. Staffell, *Renewables.ninja*, Accessed: 2026. [Online]. Available: <https://www.renewables.ninja/>.
- [43] O. Jay, R. Horadagoda, N. Bain, and M. Capon, “Reducing the health effects of hot weather and heat extremes: From personal cooling strategies to green cities”, *The Lancet*, 2019. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC6422495/>.
- [44] Ministerio de Educación, Bolivia, *Horarios escolares*, Accessed: 2026, 2023. [Online]. Available: https://www.minedu.gob.bo/index.php?option=com_content&view=article&id=6849.
- [45] P. Kapur et al., “Energy use pattern in rice milling industries – a critical appraisal”, *The Scientific World Journal*, 2012. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4571202/>.
- [46] FAO, *Rice cultivation calendar – Bolivia*, Accessed: 2026. [Online]. Available: <https://www.fao.org/4/Y4347E/y4347e0a.htm>.
- [47] Angaza, *5 solar radios with PAYG technology*, Accessed: 2026, 2019. [Online]. Available: <https://www.angaza.com/2019/02/11/5-solar-radios-payg-technology/>.
- [48] Philips, *TAR1609 – user manual, 3 W audio output*, Accessed: 2026, 2024. [Online]. Available: <https://www.documents.philips.com/assets/20240718/13ea67c9fce04b849d38b1b1008b603a.pdf>.
- [49] Instituto Nacional de Estadística (INE), *Censo de población y vivienda 2024 – cuestionario*, Question 18.8 reads *Su hogar tiene: ¿Calefón o termotanque?*. No other item of question 18 covers water heating. Accessed: 2026, 2024. [Online]. Available: <https://fm.ine.gob.bo/index.php/catalog/132/download/1404>.
- [50] S. L. Johnson et al., “Shallow groundwater temperature patterns revealed through a regional monitoring well network”, *Hydrological Processes*, 2023. DOI: 10.1002/hyp.14975.
- [51] Z. C. Johnson et al., “Continental-scale analysis of shallow and deep groundwater contributions to streams”, *Nature Communications*, 2021. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC7933412/>.
- [52] Autoridad de Fiscalización de Electricidad y Tecnología Nuclear (AETN), *Anuario estadístico de electricidad 2024*, Electricity sales by tariff category and isolated system. Accessed: 2026, La Paz, Bolivia, 2025.
- [53] Ministerio de Hidrocarburos y Energías and GIZ, “Caracterización del consumo eléctrico de los hogares bolivianos”, La Paz, Bolivia, Tech. Rep., 2022, Unit consumption values: refrigerator 458 kWh/yr, AC 451 kWh/yr, water heater 824 kWh/yr.
- [54] Fundación Bariloche and BID – Banco Interamericano de Desarrollo, “Estudio de base del sector energético boliviano”, Tech. Rep., 2018, Electric fraction of water heaters: 97.5%.
- [55] United Nations Environment Programme, *Delivering the transition to energy efficient lighting*, GEF Project ID 5299, Bolivia. Accessed: 2026, 2023. [Online]. Available: <https://www.unep.org/gef/projects/delivering-transition-energy-efficient-lighting>.
-

- [56] ELETROBRAS and PROCEL, “Avaliação do mercado de eficiência energética no brasil: Pesquisa de posse de equipamentos e hábitos de uso – ano base 2005 – classe residencial – relatório brasil”, ELETROBRAS/PROCEL, Rio de Janeiro, Tech. Rep., Jul. 2007, Brazilian household appliance ownership and usage survey, base year 2005. Used as proxy for microwave consumption in absence of Bolivian data. Accessed: 2026. [Online]. Available: https://www.procel.gov.br/sites/dev/Shared%20Documents/PPHS/Classe_Residencial_Relatorio_Brasil.pdf.
- [57] MINEM and BID, “Balance nacional de energía útil 2016”, Lima, Peru, Tech. Rep., 2016, Used as proxy for commercial end-use decomposition in absence of Bolivian data.
- [58] Ministerio de Educación de Bolivia, *RUE – registro de unidades educativas*, Active regular units (Fiscal + Convenio + Privada). Accessed: 2026. [Online]. Available: <https://rue.sie.gob.bo/accesopub/reportesrue/>.
- [59] República de Bolivia, *Ley n° 1604 de electricidad, article 48 (tasa de actualización)*, Gaceta Oficial de Bolivia, 1994.
- [60] International Renewable Energy Agency (IRENA), “Renewable power generation costs in 2024”, Abu Dhabi, Tech. Rep., 2025.
- [61] World Bank, “Bolivia rural electrification project (IDTR III) – project appraisal document”, World Bank, Washington, DC, Tech. Rep., 2023, 148 M USD, closing 2030, about 25,000 grid connections and 9,000 solar home systems. Accessed: 2026. [Online]. Available: <https://documents1.worldbank.org/curated/en/099103123132516630/pdf/B0SIB02a1785390600b80d0becc7661ef05.pdf>.
- [62] Comité Nacional de Despacho de Carga (CNDC), *Reporte de precios monómicos*, Accessed: 2026. [Online]. Available: https://www.cndc.bo/reportes/index_dos.php.
- [63] M. Hallberg and J. Hallme, “A field study about the energy demand in the rural areas of Bolivia”, Reports average daily cooking energy use in Oriente (66.7 MJ/household/day) and Yungas (68.6 MJ/household/day), M.S. thesis, KTH Royal Institute of Technology, 2015. [Online]. Available: <https://www.diva-portal.org/smash/get/diva2:838985/FULLTEXT01.pdf>.
- [64] ENDE Guaracachi S.A., *Planta solar fotovoltaica el sena*, Installed capacity 400 kWp, El Sena, Madre de Dios province, Pando. Accessed: 2026. [Online]. Available: <https://www.egsa.bo/planta/solar-fotovoltaica-el-sena/>.
- [65] C. E. Harwood and S. Almeida, “Desempenho de um aquecedor solar de água em Manaus”, *Acta Amazonica*, vol. 11, no. 2, 1981. DOI: 10.1590/1809-43921981112371.
- [66] ENERGETICA, *Proyecto el sol: Agua caliente con energía solar en bolivia*, Solar water heater dissemination programme, 2009–2012. Accessed: 2026, 2012. [Online]. Available: <https://www.energetica.org.bo/elsol/paginas/cuerpo.asp>.
- [67] Ministério das Cidades, Brasil, *Minha casa, minha vida garante aquecimento solar para 896 mil pessoas em todo o Brasil*, Published 15 July 2015. Accessed: 2026, 2015. [Online]. Available: <https://www.gov.br/mdr/pt-br/noticias/minha-casa-minha-vida-garante-aquecimento-solar-para-896-mil-pessoas-em-todo-o-brasil>.
- [68] EnergyScope, *Energyscope td documentation: Input data*, Table 29, decentralised thermal storage parameters. Accessed: 2026, 2024. [Online]. Available: <https://energyscope.readthedocs.io/en/master/sections/Input%20Data.html>.
- [69] J. Rager, “Urban energy system design from the heat perspective using mathematical programming including thermal storage”, Ph.D. dissertation, EPFL, 2015. DOI: 10.5075/EPFL-THESIS-6731.
-

- [70] ENDE – Empresa Nacional de Electricidad, *Marco de gestión ambiental y social – programa de electrificación rural con energía renovable (BO-L1222)*, SHS specs: PV $\geq$ 50 Wp, battery $\geq$ 123 Wh Li. Accessed: 2026. [Online]. Available: https://www.ende.bo/public/publicaciones/pdf/bo-l1222_marco-de-gestion-ambiental-y-social_version-final.pdf.
- [71] L. A. Panjón, “Implementación de sistemas fotovoltaicos en zonas rurales del cantón Morona Santiago”, *Revista Técnica “energía”*, vol. 6, no. 1, pp. 90–94, 2010, Overall efficiency of 0.75 for an autonomous installation supplying alternating current loads, page 91. DOI: 10.37116/revistaenergia.v6.n1.2010.233.
- [72] Lighting Global, *Requirements and guidelines for installation of off-grid solar systems for public facilities: With a focus on private-sector-supported service delivery models*, Last updated March 2023. Funded by the Lighting Africa programme of the World Bank, prepared with the Schatz Energy Research Center and ECREEE. Days of autonomy, page 25, section 5.2.7; depth of discharge for lithium batteries, page 28, section 5.2.9. Accessed: 2026, May 2021. [Online]. Available: https://www.lightingglobal.org/wp-content/uploads/2021/05/QA-InstSystFramework_2023.pdf.
- [73] A. G. García, K. Dietrich, P. C. Cutillas, C. Pérez-Andújar, and R. Palacios, “Plan nacional de electrificación de última milla en Paraguay”, Banco Interamericano de Desarrollo, División de Energía, Nota Técnica IDB-TN-3083, Jun. 2025, Table 29, page 159, medium voltage network catalogue of the Reference Electrification Model: 22.9 kV three phase line at 8,212 and 12,262 USD/km, poles included. Accessed: 2026. [Online]. Available: <https://publications.iadb.org/publications/spanish/document/Plan-nacional-de-electrificacion-de-ultima-milla-en-Paraguay.pdf>.
- [74] Instituto Nacional de Estadística (INE), *INE presenta resultados del Censo 2024*, National population 11,365,333; household size 4.1 (2001), 3.5 (2012), 3.1 (2024); electricity coverage 92.2% national and 78.4% rural. Published 28 August 2025. Accessed: 2026, Aug. 2025. [Online]. Available: <https://cpv2024.ine.gob.bo/index.php/ine-presenta-resultados-del-censo-2024-que-muestran-la-transformacion-demografica-y-avances-sociales-en-bolivia/>.
- [75] Comisión Económica para América Latina y el Caribe (CEPAL), “Panorama social de América latina y el Caribe”, CEPAL, Santiago, Chile, Tech. Rep., 2025, Observed household size across the region, from 3.0 in Chile to 4.9 in Guatemala. Accessed: 2026. [Online]. Available: <https://repositorio.cepal.org/bitstreams/df523fae-e9cd-4967-9fb3-c3dc5f7b99d6/download>.
- [76] United Nations, Department of Economic and Social Affairs, “Household size and composition around the world 2017 – data booklet”, United Nations, New York, Tech. Rep., 2017, Descriptive threshold for a small household set below 3.0, stated as a description and not a bound. Accessed: 2026. [Online]. Available: https://digitallibrary.un.org/record/3799696/files/household_size_and_composition_around_the_world_2017_data_booklet.pdf.
- [77] ENDE – Empresa Nacional de Electricidad, “Memoria anual 2021”, ENDE Corporación, Cochabamba, Bolivia, Tech. Rep., 2022, Nine distribution projects completed in 2021, totalling 763.69 km of medium voltage line and 551.17 km of low voltage line in Cochabamba, Chuquisaca, Beni and Pando. Accessed: 2026. [Online]. Available: https://www.ende.bo/public/memorias/memoria_2021.pdf.

-
- [78] ENDE – Empresa Nacional de Electricidad, “Informe de rendición pública de cuentas final 2024”, ENDE Corporación, Tech. Rep., 2025, Page 162. Medium voltage line built by ENDE DELBENI in 2024: 178 km against a target of 154 km. Accessed: 2026. [Online]. Available: <https://www.ende.bo/public/publicaciones/pdf/rpcf2024-consolidado-11.02.25-.pdf>.
- [79] ENDE – Empresa Nacional de Electricidad, “Informe de rendición pública de cuentas final 2023”, ENDE Corporación, Tech. Rep., 2024, Page 137. Medium voltage line built by ENDE DEORURO in 2023: 60.14 km urban and 38.58 km rural. Accessed: 2026. [Online]. Available: <https://www.ende.bo/public/publicaciones/pdf/informe-rendicion-publica-de-cuentas-final-ende-2023.pdf>.
- [80] Estado Plurinacional de Bolivia, *Contribución nacionalmente determinada (NDC) 2021–2030*, Submitted to the UNFCCC. Target of universal electricity access by 2030. Accessed: 2026, 2022. [Online]. Available: https://unfccc.int/sites/default/files/NDC/2022-06/NDC_Bolivia-2021-2030_UNFCCC_en.pdf.
- [81] Ministerio de Hidrocarburos y Energías (MHE), *En Bolivia, 60.000 nuevos hogares acceden anualmente al servicio de electricidad*, Average of 60,000 households newly connected to the grid per year at national level. Published 30 December 2024. Accessed: 2026, Dec. 2024. [Online]. Available: <https://www.mhe.gob.bo/2024/12/30/en-bolivia-60-000-nuevos-hogares-acceden-anualmente-al-servicio-de-electricidad/>.
- [82] S. Pfenninger and I. Staffell, “Long-term patterns of European PV output using 30 years of validated hourly reanalysis and satellite data”, *Energy*, vol. 114, pp. 1251–1265, 2016. DOI: 10.1016/j.energy.2016.08.060.
- [83] I. Staffell and S. Pfenninger, “Using bias-corrected reanalysis to simulate current and future wind power output”, *Energy*, vol. 114, pp. 1224–1239, 2016. DOI: 10.1016/j.energy.2016.08.068.
- [84] S. Bhushan et al., “Automated digital elevation model generation from multi-view satellite imagery”, *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 173, pp. 151–165, 2021.
- [85] Copernicus Climate Change Service, *ERA5 hourly data on single levels*, Retrieved November 2024. [Online]. Available: <https://cds.climate.copernicus.eu/>.
- [86] A. Magne, D. Khatiwada, and E. Cardozo, “Assessing the bioenergy potential in South America: Projections for 2050”, *Energy for Sustainable Development*, vol. 82, p. 101 535, 2024.
- [87] Ministerio de Hidrocarburos y Energías (MHE), “Balance energético nacional a nivel departamental 2021”, La Paz, Bolivia, Tech. Rep., 2022.
-