

RF Physics for DeWi Investors (Pt 1)


Networks are - by definition - at least as valuable as the sum of their individual parts. In the context of DeWi, we're particularly interested to understand: *how much is a wireless deployment worth?* If we understand the value of individual deployments, we can extrapolate a lower-bound for the value of the entire network.

This letter aims to study the links between the *physics* and *economics* of a wireless network from first principles. Part One focuses on the relationship between *costs* and *coverage*; Part Two examines the relationship between *coverage* and *value*.

First off — what is a wireless deployment?

They consist of *radios*, *antennas*, and *towers*:

- **Radios** are integrated circuits that create electromagnetic signals, which are applied to metal wires called antennas.
- **Antennas** amplify electromagnetic signals from the radio into 3-D space, enabling signals to travel over longer distances.
- **Towers** are elevated structures that house radios/antennas, enabling signals to cover a wider area.



Easy enough. Now — how much do deployments cost?

Deployments are clearly expensive in aggregate: telcos spend \$50B/yr on capex in the US alone. At a granular level, each deployment has two types of expenses. Some - like equipment, installation, and licensing fees - are paid up-

front and amortized on the income statement; others - like rent, backhaul, power, insurance, and maintenance - are paid on a recurring basis.

Let's start with the up-front costs. The largest - by far - is hardware; specifically, radios and antennas.

Radios. On most radio manufacturers' website, you'll find "datasheets" (e.g.) with information about specific radio model. Radios are differentiated by many factors, including their: power output (EIRP), power consumption, spectrum bands, channel width, peak data transfer rate, maximum number of connections, and size / weight.

Modern "carrier-grade" small-cell radios, like an Airstrand2200, costs upwards of \$50K+. These are software-defined radios (SDRs) capable of running multiple core networks and connecting thousands of end-users at once. SDRs allow developers to programmatically adjust factors like power output, frequency, and speed using nothing but software. These sorts of features were unavailable to the public just a few years ago, but are now commoditized, and prices are deflating double-digit percentages YoY as production volumes grow.

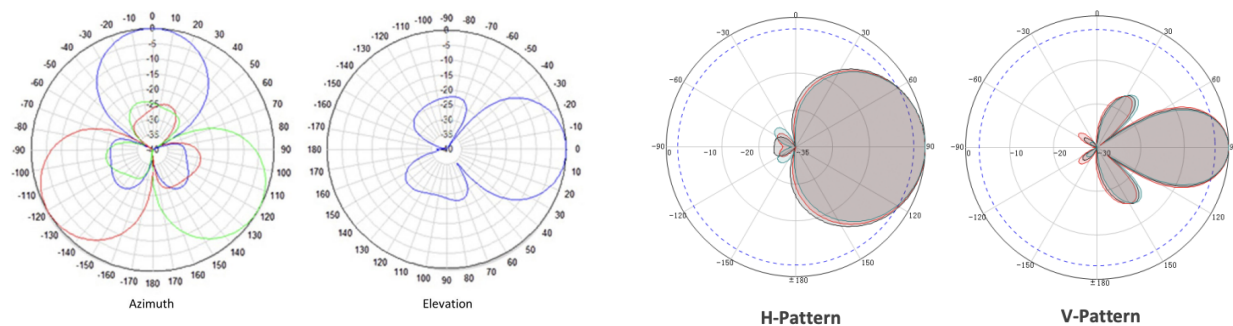
While "carrier-grade" equipment is getting cheaper, \$50K is still too expensive for most. Instead, DeWi networks have chosen to focus on inexpensive mass-market hardware, allowing for a broaden reach and enabling as many people as possible to participate in the growth of the network — an admirable goal. To be more specific, the table below benchmarks the six radios currently sold on Pollen's website: prices vary between \$750-\$7.5K, power outputs between 30-53dBm (the FCC limit for CBRS is currently 47dBm), maximum of 32-288 concurrent connections per radio, and max speeds of 110-330Mbps down (14-68Mbps up).

Pollen Mobile: Radio/Antenna Overview							
		Nova 436Q	Nova 227	X-300i	FW-300i	Flower	Bonzai
Manufacturer		Baicells	Baicells	Blinq	Blinq	Moso	Moso
MSRP		\$5,299	\$749	\$4,299	\$7,499	\$1,599	\$999
Radio Dimensions							
Max Carriers	#	2	1	3	3	2	2
Concurrent Users	#	192	32	288	288	?	?
Power Consumption	Watts	60	18	70	48	21	21
Power Output (EIRP)	dBm	30	27+13	33	36+17	45	30
Peak Rate - Up	Mbps	68	14	33	33	28	28
Peak Rate - Down	Mbps	290	110	330	330	220	220
Max Range	km	12	5	—	—	—	—
Latency	ms	30	30	—	—	—	—
Channel Width	MHz	5/10/15/20	5/10/15/20	10/20	10/20	—	—
Supported Bands							
Band 48 - US CBRS	3.55-3.70GHz	Yes	Yes	Yes	Yes	Yes	Yes
Band 42 - EU CBRS	3.40-3.60GHz	Partial	Yes	Yes	Yes	No	No
Band 43 - C-Band	3.60-3.80GHz	Partial	Yes	Yes	Yes	No	No
Physical Dimensions							
Height	cm	33	25	33	35	41	27
Length	cm	24	25	23	57	29	21
Width	cm	11	8	21	23	12	8
Weight	kg	8	3	5	12	6	2

Note: all columns include radios + antennas, except for the Nova436Q (radio only).

Telcos already provide 25-50Mbps connectivity reliably using licensed low/mid-band spectrum. Reflecting on the table above, a DeWi pessimist might say that the current deployments have limited utility, given that the radios would at best improve speeds from 25-50Mbps to 110-330Mbps while also requiring significant investment in densification. We note this pessimism would apply equally to Helium, Pollen, and XNET, all of whom work with more or less the same set of hardware manufacturers (at least for now).

Antennas. Radios can power a single or multiple antennas that amplify signals into 3-D space. Antenna design is an entire (and entirely fascinating) field of study, but the concepts are simple. You can think of an antenna as a vector field that amplifies radio waves into 3-D space. Intuitively, there is a tradeoff between coverage and capacity: antennas that are “focused” on a narrow area can provide more capacity, *ceteris paribus*, than an antenna focused over a wide areas. As an analogy, you can imagine the difference between the light emitted by a candle vs a laser pointer. *Beamwidth* is the technical term for the aperture (angle) along which the majority of an antenna’s power is concentrated.



There is, of course, the question of absolute amplification power regardless of direction — in other words, the vector magnitude. Manufacturers generally disclose a figure called *antenna gain*, measured in dBi, that represents the logarithm of the antenna’s maximum amplification factor along its most powerful axis. An antenna with +5dBi gain amplifies signals by a factor of 3x [$3x = 10^{(5/10)}$] along its most powerful axis; one with +17dBi gain amplifies by a factor of 50x [$50x = 10^{(17/10)}$].

In practice, antennas help mitigate the tradeoffs of moving into higher frequencies. We’ll cover this further, but for now understand that to transfer more data, antennas must transmit over higher frequencies, which means signals degrade much faster over distance. Antennas that are highly “focused” can offset this signal degradation and still provide high-capacity connectivity over long distances. This is the fundamental concept behind point-to-point (PtP) mesh technology, which uses highly-directional antennas (hence the name *point-to-point*) transmitting over high-frequency spectrum (60-80GHz) to provide an extremely fast internet connection (10Gbps) over extremely long distances (10km). At the other extreme are WiFi routers, which transmit in all directions, but suffer from severely limited range.

Other Up-Front Costs. Other equipment, such as mounts, wires, and casings, add costs but are typically not meaningful enough to move the needle on deployment economics. Licensing costs are highly specific to individual deployments - for example, most CBRS deployments require only a nominal SAS registration fee (<\$25), while a new full-scale tower in a suburban NIMBY neighborhood might pay >\$25K in municipal licensing and related legal fees - so we ignore them for our purposes of a generalized valuation framework. Installation costs are similarly variable, although we typically think of 2-4 hours of technician time as a helpful benchmark for the labor costs of a single small cell deployment.

Now, on to the recurring costs of wireless deployments: rent, backhaul, and power.

Rent. In TradWi, telcos previously owned their tower infrastructure but eventually sold them and must now pay ongoing rents to landlords, namely tower REITs. In DeWi, deployments exist on a spectrum, from having zero rent

(consumer deployments), to having token rents (hosted model), to having traditional fiat rents. In the consumer model, end-users install the radio in their home or apartment and incur zero incremental rent. In the hosted model (see [Fairspot](#) or [HeliumDeploy](#)), hosts earn a 20-70% share of rewards, without contributing a single dollar of capital. In the traditional model, deployments compete directly with traditional operators and pay the same fiat-denominated rents as everybody else.

Rent prices vary based on many different factors — as an example, a tall building in a major city might charge \$[100]/mo for roof space, whereas an actual tower starts at \$[1000]/mo. In addition to rent, landlords often charge additional fees every time an operator touches the tower, such as for routine maintenance or upgrades. For dedicated radio towers, months-long negotiation process, multi-year contracts, and annual price increases are the norm. In the US, pricing increases are typically fixed; emerging markets often peg to inflation.

Backhaul. “Backhaul” simply means a connection to the Internet, typically via a series of wires and/or radio signals to one of the 126 [internet exchange points](#) in North America. In practice, a single high-capacity backhaul connection might cost \$1-2K/mo and allow for 10Gbps of internet capacity (10,000Mbps). Given the average household uses 2-4Mbps during peak hours, the single backhaul connection can be resold to 2.5-5.0K households bringing per-household costs down to <\$0.50/household/mo at full utilization. Naturally, this means that users’ internet speeds are capped during peak hours when utilization is high (and explains why cable providers are buying youtube and instagram [ads](#) trashing telco fixed wireless’ offerings).

On the other hand, if less than 100 households are sharing a connection (backhaul costs of >\$10-20/user/mo), the gross margins on providing internet access can be deeply negative, as many bankrupt WISPs can attest to (see: [Starry](#)). Note that \$1-2K/mo represents the high end of the range... a smaller deployment on an apartment building with existing fiber may pay as little as \$50-100/mo.

Power. Radios obviously need electricity to transmit electromagnetic signals. The amount of power required depends on a radios average power consumption. The current outdoor DeWi radios are rated for 50-60 watt power consumption, meaning power costs are modest relative to other costs. At the US average of \$0.15 per kilowatt-hour, outdoor DeWi radios use only \$5-6/mo worth of electricity. In 2021, telcos consumed an estimated 470TWh of [electricity](#) globally, suggesting a (very rough) estimate of \$70B+ annually to power the world’s telcos networks.

Other. Maintenance, insurance, and other administrative costs can also be significant ongoing expenditures depending on the nature of the deployment.

Now we have a framework for how much a deployment can cost. Before determining how much coverage is *worth*, we first need to know how much coverage *exists*.

We like the term useful coverage, which is an output of three factors:

- How powerful is the equipment?
- What spectrum is being used?
- How high is the antenna, relative to its surroundings?

How powerful is the equipment?

Radios are designed to operate at a maximum *power*, measured in watts. One watt is equivalent to the power required to accelerate a 1kg mass at 1 m/s² over a physical distance of 1m and temporal distance of 1s. The nature of radio waves means engineers often need to deal with power levels ranging from 100W, to 0.01W, to 0.000001W — since trailing zeros are annoying, engineers use a logarithmic measure, *dBm*, where 0dBm is defined as 1/1,000th of a watt. 10dBm is ten times more powerful than 0dBm (1/100th of a watt), 30dBm is one hundred times more powerful than 10dBm (i.e. 1 watt), and so on. Here's a simple [tool](#) for converting between the two.



Power limits are regulated because excessive power can cause interference for nearby users, particularly in lower frequencies. In CBRS spectrum, the FCC allows two categories devices: Class A, with a power limit of 30 dBm (1 watt), and Class B, with a power limit of 47 dBm (50 watts). The licensed bands directly adjacent to CBRS, which therefore have similar physical propagation characteristics, are allowed power limits hundreds of times higher (72-75 dBm).

Some telcos have lobbied the FCC for higher power limits in CBRS: notably AT&T in '19 asked 62dBm limits (1.6 kilowatts) and Dish in '21 asked for carrier-grade 72dBm limits (16 kilowatts). Their efforts saw some political support, but were opposed by T-Mobile - who already owns nearby 2.5GHz spectrum and would rather handicap competitors - and by WISPs - who want the spectrum for their fixed wireless offerings. The FCC has yet to formally approve power limit increases in the CBRS band, but we're optimistic they will given the success of the band.

For further reading on the current state of CBRS, see recent opposing letters published by the [CTIA](#) (telco industry association) and by a [CBRS industry consortium](#).

Radios are only part of the story. The antennas currently in use by the DeWi networks offer gains of up to +3dBi for indoor and up to +17dbi for outdoor. Note that FCC power limits apply *after* antenna gain, so the most powerful legal CBRS setup today is a 30dBm radio with +17dBi antenna (30+17=47dBm limit). In other words, a 1 watt radio amplified up to 50 watts.

Receivers also have a *gain*, i.e. a factor by which they can amplify the electromagnetic signals. Mobile phones have receiver gains of up to 3dBi, effectively doubling the power of received signals. The FCC sets power limits for CBRS receivers as well as transmitters at +23dBi. This limit applies to devices such as Helium/Pollen mappers.

Adding *radioPower(30dBm) + antennaGain(17dB) + receiverGain(3dB)* results in the theoretical maximum power of a deployment *if one is standing directly next to the antenna along its most powerful axis*. In this case, with a CBRS small cell transmitting to a mobile phone, maximum received power is 50dBm, or 100 watts. Licensed deployments on adjacent frequencies are allowed to operate at 75dBm, or 31,000 watts... 300x more powerful than CBRS.

What spectrum is being used?

Radios send electromagnetic signals at a certain frequency, measured in Hz — one Hz is equivalent to 1 cycle per second. Human voices resonates at 100-200 Hz, while radio waves span from 10KHz to 300GHz. Some bands in the radio spectrum, like CBRS (3.55-3.70GHz) and WiFi (2.40-2.48 & 5.15-5.85GHz), are approved by regulators for general public use. Others, like C-Band (3.90–3.98GHz) and mmWave (24-27GHz+), are auctioned off exclusively to licensed carriers. In the end, both licensed and unlicensed bands are subject to FCC power limits and other forms of regulation. To state the obvious, there's no *physical* difference between unlicensed vs licensed bands - the laws of


physics apply regardless of government status - but there is a critical difference between high- vs low-frequency spectrum.



Since speed of light is constant, if we know a wave's frequency we must also know its *wavelength*, or the distance traveled by the wave in a single cycle. Divide 300 (i.e. the speed of light in millions of meters per second) by the frequency in question measured in MHz (i.e. millions of cycles per second) and the result is the wavelength measured in meters. To give some intuition, low-band cellular spectrum (900MHz) has a wavelength of ~1/3rd of a meter, CBRS spectrum (3.5GHz) has a wavelength of ~1/12th of a meter, and mmWave spectrum (27GHz) has a wavelength of ~1/100th of a meter. Here's a simple tool for converting between the two.

Lower frequencies (longer wavelengths) propagate further and better resist interference, but have limited data transfer capacity. Higher frequencies (shorter wavelengths) have high data transfer capacity, but propagate over short distances and are far more susceptible to interference.

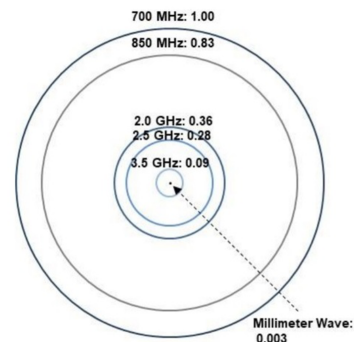
Because propagation characteristics vary fundamentally across different bands, telcos take a portfolio approach to spectrum licenses: they use lower frequencies to provide broad coverage with powerful macro radios. Then, after the base layer of the network is built, they densify networks with smaller radios that transmit on higher frequencies to enhance data transfer capacity. Already today, the vast majority of *coverage* is over 700-900MHz, but the vast majority of *data transfer* is over 1.7-2.5GHz.



Source: Tutela

Relative Propagation Distance

700-800MHz vs 2.0-2.5GHz vs 3.5GHz vs 24GHz



Source: MoffettNathanson

Appetite for “high-frequency” spectrum has changed predictably, but dramatically, over time. In the 1920's, AM radio stations transmitted on <1MHz, with wavelength the size of a football field (in fact, there are stories of early radio stations in Chicago interfering with stations in Tokyo). In the 1930s-1950s, FM radio stations transmitted on 50-100MHz with wavelengths of roughly a meter. In the 2000's, during the early days of the cellular industry, “mid-band” was 1-2GHz... today, frequencies as high as 5-6GHz are considered “mid-band”, given that telcos are investing into spectrum in the 28-39GHz range (aptly named *millimeter-wave*) and even higher. As pictured above, the propagation distance of a mmWave signal is minuscule relative even to CBRS spectrum, let alone any lower-frequency licensed spectrum.

Despite its faults, higher-frequency spectrum is attractive primarily because it is so underutilized. Given sub-6GHz spectrum has been usable for decades, even the tiniest specks of spectrum (say, a single 10MHz channel) are highly contested. Any application that needs multiple contiguous 10MHz blocks is simply out of luck. On the other hand, in the 24GHz, 60GHz, and especially 80+ GHz ranges, it's easy to find large contiguous chunks of tens or even hundreds of MHz of available spectrum. This allows for a higher throughput of data transfer, i.e. faster internet speeds. Crypto-native folks may find the analogy to block widths instructive: bigger blocks (or higher frequencies) allow for higher throughput, but come at the cost of higher requirements for block producers (or network operators).

Shifts in spectrum utilization happen gradually... and then suddenly. Millimeter-wave spectrum, considered worthless a decade ago, is now a multi-billion dollar asset on telco balance sheets and is actively being used to drive carrier capacity in dense areas. For example, Verizon has deployed 30K+ mmWave small cells across the US: in some cities (Chicago), already half of all customer data travels over mmWave spectrum, and in others (San Diego), the average observed download speeds are in excess of 1,000Mbps!




Table 1: San Diego 5G UWB Average DL/UL Speeds Across Areas

Map Area	Neighborhood	Downlink Speed (Mbps)	Average Uplink Speed Average (Mbps)
D	City Heights	1,207	48
E	College	1,883	91
H	Downtown	1,088	65
E	Eastern	NA	NA
D	Kensington	1,223	53
A	Linda Vista	1,344	56
B	Mission Valley	1,055	55
C	Normal Heights	NA	NA
C	North Park	895	36
G	Uptown	1,232	54
Total		1,241	57

Source: ETL Wireless Research LLC

High-frequency spectrum is wonderfully abundant, but exacerbates the fundamental problem in the wireless business: distance. Recall that the power of an electromagnetic signal decays exponentially not only with respect to frequency, but also with respect to *distance*. The combined effect is called *path loss*, and is expressed in logarithmic terms (dB), where:

$$-PathLoss = 20 * \log(df) + 32.44$$


with distance d in kilometers and frequency f in MHz.

For example, a low-frequency (900MHz) cellular signal traveling the length of a football field (0.1km) would see a path loss of -72 dB [$20 * \log(900 * 1/10) + 32.4$]. Remember dB are logarithmic units, so -72dB means losing roughly eight orders of magnitude of signal strength (a >99.99999% loss). Moving up in frequency to CBRS spectrum, path loss over the length of a football field is -84dB [$20 * \log(3600 * 1/10) + 32.4$]. In other words, CBRS is more than order of magnitude worse; to cover a similar area, one must deploy >10x as many radios. For spectrum above 11GHz, propagation is *another* order of magnitude worse; covering a similar area requires >100x as many radios. Such are the brutal economics of building wireless networks.

Decibal Loss (dBm)										
				Frequency (MHz)						
Distance	ft	mi	km	900	1,800	3,600	7,200	14,400	28,800	57,600
	33	0.0	0.01	-52	-58	-64	-70	-76	-82	-88
	164	0.0	0.05	-66	-72	-78	-84	-90	-96	-102
	328	0.1	0.10	-72	-78	-84	-90	-96	-102	-108
	820	0.2	0.25	-79	-86	-92	-98	-104	-110	-116
	1,640	0.3	0.50	-86	-92	-98	-104	-110	-116	-122
	3,281	0.6	1.00	-92	-98	-104	-110	-116	-122	-128
	8,202	1.6	2.50	-99	-106	-112	-118	-124	-130	-136
	16,404	3.1	5.00	-106	-112	-118	-124	-130	-136	-142

Note: highlighted column represents CBRS.

In addition to path loss, higher-frequency spectrum also exacerbates the impacts of *interference* and *attenuation*. Interference occurs when two electromagnetic signals travel near each other, and attenuation occurs when an electromagnetic signals passes through physical matter. At higher frequencies, the impacts of interference and attenuation become exponentially worse with distance: on 10GHz spectrum, in heavy rain, atmospheric signal attenuation is 1dB per km (i.e., a 26% drop in signal strength every kilometer), but 3GHz sees barely any attenuation (0.01db, or a 0.2% drop, per kilometer). Similarly, In 2.4GHz WiFi spectrum, traveling through a concrete wall cuts signal strength down by -10dB (a 90% drop), while the same wall on 5GHz WiFi might cut signal strength -20dB (a 99% drop). This explains why you may find yourself inside a brick/concrete buildings able to make a phone call on 900MHz, yet unable to stream music on 2.5GHz.



Source: everythingrf.com

How high is the antenna, relative to its surroundings?

The higher the tower, the more of its surroundings the antenna can “see” without interference (direct line-of-sight). Note that we use the term *tower* generically, since any structure with sufficient elevation relative to its surroundings - e.g. apartment buildings, billboards, water towers, or even street lamps - can also work. Towers tend to be 50-400 feet tall, with rent prices increasing the further up the tower the equipment is placed.

From first principles, say we have a building of certain height and a nearby tower that’s a given multiple taller than the house. Then, for every unit of distance between the house and the tower, the radius of blocked coverage (the “shadow” caused by the building) is equal to the inverse of the tower:building ratio.



To use a real example: the average two-story home is 20-feet tall. Let's assume there's a 100-foot tower (5x taller than the building) located 1 mile away from the tower. Then, directly next to the house, there is no direct line-of-sight (a coverage "shadow") for the length of 1/5th of a mile, or 3.5 football fields. If the tower were instead 400-feet tall, the shadow would instead be only 1/20th of a mile, or 1 football field. This rule follows from basic trig, if you ignore curvature of the earth — we'll leave the proof as an exercise to the reader.

Note that it's the *relative* height of the tower and the building that matters. In Manhattan, with a hundred skyscrapers taller than 500 feet, it would be physically impossible to build a tower tall enough to provide coverage in the gaps between buildings (real-world solution: hang radios off the side of buildings). On the other hand, a 400-foot tower in the midwest might provide line-of-sight coverage up to 10 miles away, assuming it can actually transmit that far. Remember that even with direct line-of-sight, the power, distance, and frequency must still jive — a mmWave deployment cannot reach 10+ miles away, even with direct line-of-sight, simply because the path loss over such a long distance would be insurmountable at such high frequencies (>150dBm).

Implications for CBRS

The majority of mobile data traffic today happens over 1.7-2.5GHz, which is meaningfully different than CBRS (3.55-3.70GHz). First, FCC limits power output in CBRS bands at 47dBm, 16dB lower than in the 2.3-2.4GHz band (63dBm). Second, by nature of having a shorter wavelength, CBRS transmissions have an incremental 3-5dB of path loss relative to 2.3-2.4GHz. These two factors add up to a ~20dBm spread between CBRS vs licensed mid-band spectrum, before counting for interference or attenuation (which are exacerbated at higher frequencies).

Empirical results support these ranges. For example, studies by [OpenSignal](#) show a ~15dBm spread between Verizon's 3.5GHz and AT&T/T-Mobile's 2.3-2.5GHz performance. And while range is significantly shorter in CBRS, the results clearly show the advantages of the band: download speeds more than twice as fast as 2.3-2.5GHz.



In summary, we hope this letter gets a few points across:

- Higher-frequency spectrum, including CBRS, is the only path to providing 100+ Mbps internet for the masses. Leveraging this spectrum requires densifying networks by at least an order of magnitude... a feat that is impossible under the legacy telco model.
- CBRS spectrum can be massively valuable, and is already being used to serve customers at scale. However, the true potential of the band is limited by FCC power limits that are 300x lower than in licensed bands. The changing of this rule will be a huge step forward for DeWi.
- On a 3-5 year timeframe, the current batch of DeWi radios will be more or less obsolete. The equipment currently available at retail-like prices is simply too limited in capabilities - in terms of max # of connections, max capacity,

max power - to be broadly useful the way many folks in the DeWi community are expecting.

Part Two of this letter will attempt to answer: *how much is coverage worth?*

It turns out to be a complex question with a number of factors that often influence each other:

- How many users are covered by the deployment? (density)
- How much data do they use? (demographics)
- What percentage of traffic can be captured? (competition)
- What's price are users willing to pay? (elasticity)
- What currency is used to pay for data transfer? (inflation)

As always, we welcome feedback from our partners, friends, and skeptics.

Until next time 🚀