RMI Insight Brief

Drilling Down on Gas Loss

Supplement: Methodology

A new framework for correcting self-reported oil and gas data

Dwayne Purvis (consultant), Kevin Gauthier, Deborah Gordon, Carmela Chaney, Lauren Schmeisser, Cayla Calderwood, Mathew Sofield (consultant)
September 15, 2026

Table of Contents

1. INTRODUCING THE FRAMEWORK — BOTTOM-UP SYSTEMS METHOD 1
2. DETAILED METHODOLOGY 2

Data Available from Regulator 3
Data Collection and Reduction 5
Data Normalization 8
Analysis and Interpretation of Reported Disposition 9
Extension to Actual Gas Volumes Released 19
Step-by-step for Inference of Actual Volumes 23
Validation of methodology, internal controls 25
Validation of methodology, external controls 26

3. FRAMEWORK LIMITATIONS AND FUTURE DEVELOPMENT. 34

Limitations 34
Future development opportunities 34

4. ACRONYMS AND ABBREVIATIONS 35
5. KEY FORMULAS AND CONVERSIONS 36

5.1 Gas-to-Oil Ratio (GOR) 36
5.2 Percentage of Gas Wasted 37
5.3 Inferred Gas Production 37

ENDNOTES 38

1. Introducing the framework — bottom-up systems method

Upstream oil and gas production emissions, particularly of self-reported uncombusted methane, have proved difficult to verify. Bottom-up measurements and self-reported estimates by companies and agencies have consistently reported methane emissions far below the methane measured empirically in the atmosphere by aerial and satellite surveys. Although empirical measurements are more accurate, they cover limited areas over select periods of time.

The *Texas Wasted Gas Analysis* attempts to bridge the gap between self-reported estimates and empirical observations by quality-testing and correcting the volumes of gas produced and disposed of, and to extend that methodology across the entire state.

The concept of forensic analysis and correction of large-scale data sets is not new, but it has not historically been applied to oil and gas. Bridging the concept to oil and gas required the creation of new tests and corrections that are unique to this environment and oil and gas production data.

Self-reported data is routinely analyzed across sectors, scrutinized, and corrected to infer operational realities. Osiecki et al. (2024), for example, present a recent forensic analysis of self-reported Toxic Release Inventory data. Another well-established analogue can be found in occupational health surveillance, where analysis of self-reported data has consistently exposed underreported injuries, illnesses, and contributing factors (Kyung et al., 2023).

Methodological analogs also suggest that the anomalies observed in Texas self-reported oil and gas data are not unique. A common failure mode in forensic analysis of self-reported data across sectors is the treatment of missing data, particularly when it is “Missing Not at Random” (MNAR), as observed in this accounting analysis. Zhou, Aryal, and Bouadjenek (2024) show that ignoring or incorrectly imputing MNAR data can lead to systematic underestimation of the severity of a condition, and that addressing MNAR typically requires explicit modeling of the missing data mechanism itself.

When companies are required to report production and disposition of oil and gas, the implicit assumption is that all stakeholders report the data and report it accurately. This work demonstrates that missing or unreliable data should not be interpreted as zero, and how it can be replaced with more reliable data.

As no precedent was found in the literature, this methodology was developed from the ground up for application in Texas by Purvis Energy Advisors in partnership with RMI. It maintains granularity at the lease level and accuracy by using existing data that comes directly from the operating companies. It incorporates structured internal checks to identify reporting anomalies and applies external consistency tests to validate results using independent information. Together, these elements strengthen confidence in modeled outcomes and provide a robust foundation for policy design, targeted inspection, and climate-aligned decision-making.

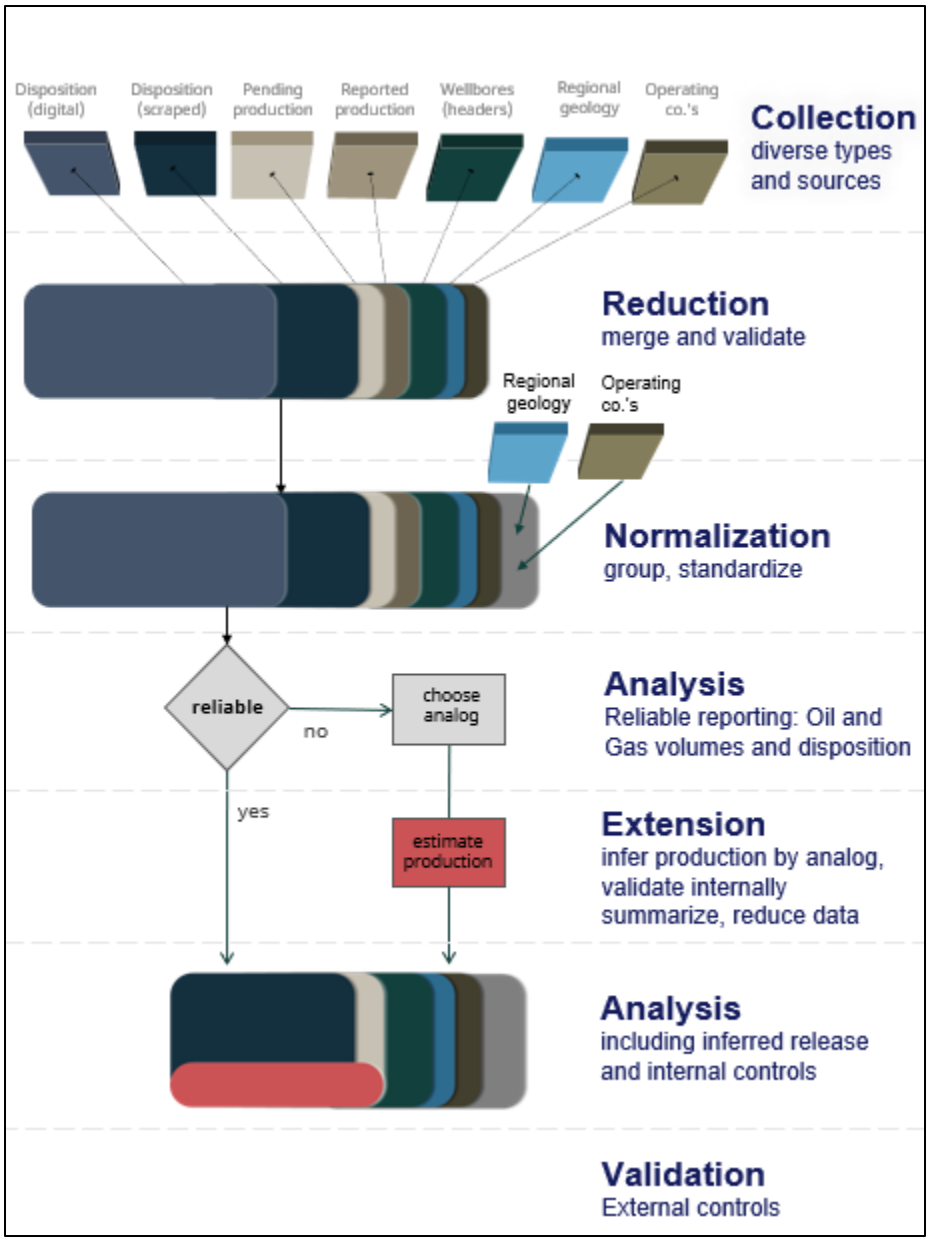
2. Detailed Methodology

Exhibit 1 below shows our methodology in a simplified schematic. The methodology begins by collecting and merging multiple datasets from the state regulator of oil and gas, especially data on the disposition of produced volumes. To facilitate initial analysis, we normalized key data fields and made simple interpretative calculations. The resulting information is then analyzed for quality, testing specifically for several ways that the data may be unreliable. In the end, we parse the records into two categories: unreliable to be replaced and reliable to be used as the basis for replacement.

In the next step, we choose a set of analogous leases with reliable data based on two dimensions—similar oil production rates and similar geology—selecting the nearest analog group with a minimum number of samples. We iterated to find the best methods for identifying analogs, and in the end, we ran an alternative scenario for the choice of leases to include among the analogs. We tested our algorithm by comparing the model estimates against reported gas production for reliable records.

Although the original data included only one month in each record, we aggregated the production in each lease to create one record for each year. We made the calculations of corrected gas production on these one-year summary records, and we tested our calculations against top-down aerial and satellite measurements of methane in the atmosphere. Finally, we analyzed and visualized the results.

Exhibit 1. Schematic flow of the methodology.



Data Available from Regulator

Most information about oil and gas operations remains confidentially within oil companies, away from competing and conflicting interests, but some of the most important information becomes public knowledge as it is reported to regulatory agencies. Regulatory data is the most common, most direct, and most readily available for detailed analysis. Regulatory data often includes information on, for example, well locations, well construction, and production among other data points.

The methodology relies almost exclusively on public data available from the Railroad Commission of Texas (RRC) which, despite its archaic name, regulates oil and gas production in the state. We also relied on regional structural interpretations made by the Bureau of Economic Geology which functions as the geologic survey for the state and county outlines from Texas State Department of Transportation. From these sources, we collected, combined and interpreted diverse tables of data delivered separately including information about wells drilled in the state, what volumes were produced, and how they were disposed of.

Disposition data via RRC web query tool

Although unavailable in bulk digital downloads, the new disposition information was accessible through an interactive, form-based query tool hosted on the RRC website.[i] The tool returns results for a single Lease ID per request and displays the output as a structured HTML page. The RRC does not provide an application programming interface (API) or bulk download mechanism for this data.

To retrieve the information at scale, we developed an automated data collection process using Python to replicate user interactions with the web interface. The procedure consisted of the following steps:

1. Submitting form-based requests for each Lease ID in the state, using a comprehensive Lease ID list compiled for other disposition tables.
2. Replicating the session cookie required for the results page to load successfully.
3. Parsing the structured HTML response returned for each Lease ID.
4. Extracting disposition fields and writing the results to a flat ASCII data file.

To comply with service constraints imposed by the RRC, requests were throttled to the maximum allowable frequency. Even with rate limiting, the process was subject to intermittent failures caused by unexpected data values or server-side responses. As a result, the collection process ran slowly and required slightly more than three weeks to complete.

Shortly after completion, the RRC modified the behavior of its online query tool in a way that disabled the scraping algorithm used in this process. While a revised approach may restore functionality, continued access via this method cannot be assumed and remains dependent on externally maintained web infrastructure.

Several limitations affect the completeness and reliability of the disposition data. In some cases, operators are permitted to estimate released volumes rather than measure them directly. Additionally, as of late 2024, a substantial number of operators continue to fail to report under the revised standards. As with most public reporting systems, data for the most recent reporting month is routinely incomplete due to late submissions. Accordingly, the most recent month of available disposition data has been treated as incomplete.

The accuracy of the scraped dataset was evaluated through multiple validation steps. Individual Lease ID results were manually checked against the RRC web interface to confirm correct parsing. Aggregated volumes were compared across internal totals derived from the scraping process and

[i] Specifically, the Production Reports Query (Form PR) found on this page: https://www.rrc.texas.gov/resource-center/research/research-queries/.

against total released volumes reported in the RRC's digital disposition dataset. These checks confirmed internal consistency and alignment with published totals.

Combined, the disposition data serves as the base table to which other datasets were merged or referenced for analytical purposes. Long-term trend analysis relies on the standard digital disposition data, while analyses requiring the enhanced reporting detail introduced in 2022 rely on the web-query–derived dataset. Additional data fields were merged using district and Lease ID as join keys. The example data set and data key demonstrate the end result of this process.

Data Collection and Reduction

Analysis required the simultaneous use of multiple kinds of data reported separately by the RRC. Following are the sources and types of data which were loaded into an SQL database then merged into combined tables for the purpose of analysis in TIBCO Spotfire (TM).

The agency makes data available as flat files for download and through web-based query tools that generate smaller datasets, so it comes as individual tables or individual records. The table of wellbores in the state contains one record per well including data fields pertaining to static information about the well. Tables of production and disposition include less identifying information but many more records. Each record represents a combination of reporting entity (in this case, RRC lease numbers a.k.a. Lease ID) and the month of activity. We combined the tables using two keys: the combination of the RRC district number and RRC lease number (because the numbers can be repeated in different districts), and the month of activity.

Disposition data via bulk download

Disposition of produced volumes—namely, where the fluids go after having been produced—is reported to the RRC separately from the total volumes produced and in greater detail. Disposition information, for example, quantifies how much was sold, spilled, stolen or used on location. The data has been available in digital form by reporting entity, namely RRC Lease ID, [ii]since 1993. Disposition data is the core of the analysis because it includes operators' self-reports of venting and flaring and has been the basis of statewide reporting of emissions.

The data included two kinds of duplication that needed to be removed. First, there were multiple entries in the rare situation of a single well producing from two separate zones with separate RRC Lease IDs. In this case, we combined the multiple producing sections of the well and grouped them as one entry in our database. Second, we found some duplicates when corrections had been filed, and we kept only the corrected data.

We compared the results with other sources of the data, including analyses or summaries published by the RRC, to confirm our results.

A change to RRC reporting rules in early 2021 required operators to provide additional detail on the disposition of released gas beginning with January 2022 production. Under the revised rules, venting and flaring must be reported separately rather than as a combined category, and operators must

[ii] The RRC offers information on two web pages. This link contains data available for download: https://rrc.texas.gov/resource-center/research/data-sets-available-for-download/. And this link contains tools to query RRC data and return limited results: https://www.rrc.texas.gov/resource-center/research/research-queries/.

identify the regulatory authority under which the release occurred. These changes were expected to represent a step-change in the quality and interpretability of publicly available disposition data.

In practice, two factors delayed the realization of this improvement. First, operator compliance increased only gradually following implementation, reaching a maximum—though still incomplete—level in early 2023. Second, the RRC's legacy data systems and bulk download products were not updated to accommodate the new reporting fields. As a result, the additional disposition detail did not appear in the bulk download of digital disposition information.

Open-records request of disposition data for redundancy and reproducibility

In parallel with the web-based data collection, the same disposition data were requested through a formal Texas open-records request.[iii] The RRC accepts such requests by email, and effective requests must state the desired data clearly and precisely. In this case, the requested data was delivered without advance notice, approximately two and a half months after submission.

By the time the open-records data was received, the web-query dataset had already been compiled and validated, making the open-records delivery operationally redundant. Nevertheless, this pathway is documented here to support methodological reproducibility. In jurisdictions where automated web access is infeasible or prohibited, an open-records request may represent the only viable mechanism for obtaining equivalent data.

Production data via bulk download

We also downloaded and examined the digital record of reported production volume from the RRC. This is a foundational type of data used for quantitative analysis, including only the volumes that came out of the well(s) in each month. The disposition data pertains to the month in which the volumes were disposed of, which can be slightly different in the case of oil, which can be stored in tanks. The disposition data also, as described above, breaks down the volumes into subtotals by destination such as sales versus released. This data has some of the same idiosyncrasies as the disposition data, but we merged only one record from the production table for each of the normalized records in the disposition table. Specifically, the production data table provided important header information about the leases — such as field, reservoir, and district — as well as validation of the total volumes disposed of.

Pending production data collection

Digital data is also available for wells not yet assigned RRC Lease ID numbers, but in a separate table called "pending production." Instead of using an (unassigned) RRC Lease number, the records are assigned a unique temporary identifying number. Because we were conducting a single snapshot analysis, we included the pending production data using the temporary identifier in place of the RRC Lease ID.[iv]

This data was downloaded and merged into the table of disposition data in the state, adding only about 3% of production and not changing significantly the proportions of flaring and venting.

[iii] For more information, visit https://www.rrc.texas.gov/general-counsel/open-records/procedures-for-requesting-information/.

[iv] On the RRC page "Data Sets Available for Download," under Production Data is "Production Report for Pending Leases" https://rrc.texas.gov/resource-center/research/data-sets-available-for-download/. Each release is a continuation, so it's only necessary to download one file.

Well headers data collection

We downloaded the RRC's list of oil and gas wellbores for a single, contemporaneous point in time, and we merged header data into the summary table including the number of plugged and unplugged wells in each lease, completion dates (first, average, last), the average latitude and longitude of each lease, and the county in which the lease is located. We did have to aggregate the data from individual wells to lease level to merge with the disposition data, and we used this info to map and characterize the leases responsible for the emissions.

Operating companies data collection

We downloaded the contemporaneous directory of companies in the state with licenses to operate issued by the RRC, including for production as well as transportation and other activities. More specifically, these are more than 7,600 companies with valid or recently expired P-5 forms on file, including company name, unique identifier, address and contact information. As described below, we used this information to identify companies within common control so that we could summarize emissions by companies, including corporate subsidiaries.

It should be noted that we did not try to merge information from the Secretary of State about e.g., corporate officers, which could be used to help define corporate families.

Regional geologic structures data collection

For interpretation of regional geologic (structural) areas, we integrated several maps of structural features to find the leases within each. The Bureau of Economic Geology's (BEG) published shapefiles with its Permian Basin GIS Database.[1] For the rest of the state, we digitized major structural elements from the BEG's Atlases of Major Texas Oil and Gas Reservoirs,[2] but these sources are so old that we incorporated the updated interpretations from Hammes and Gale (2013) as necessary. Finally, these major structural elements shapefiles were reconciled to remove gaps and overlap.

To locate the leases, we averaged the available latitude and longitude coordinates of the leases' individual wells, and we posted the lease locations in ArcGIS along with the structural features. Finally, we executed an intersection between the layers—geologic areas and average lease locations—to assign each lease to a structure.

Finally, for visualization, we also collected and loaded GIS shape files for counties from the Texas State Department of Transportation Open GIS download site.[3]

Data not collected and excluded

We reviewed some other data that we thought might be of value, but in the end the following were not included in our analysis.

Applications for venting and flaring: A separate online query tool contains data about applications for permission to vent and flare, but we did not find this data valuable for this work.

Completion information: Similarly, we considered whether regulatory information about the initial completion and testing of wells might contain information or implications about releases. However, the data available digitally proved unhelpful. It might be possible to extract useful information from parts of the submission form that are not already digitized.

Data Normalization

In some cases, the range and variety of values were quite broad, so we grouped them into fewer, more similar groups and iterated the definitions of those groups as necessary to improve the clarity of the trends.

Types of disposition binning

We simplified the nine categories of disposition into subtotals, and we characterized the nature of the disposition reported. For example, we combined disposition categories of "gas lift" and "field ops" into "used on site," and we combined "repressure" and "pressure maintenance" into "injected." Then we binned each entry according to the nature of reporting of gas: none vented or flared, some vented or flared, all vented or flared, and no gas reported with oil production. The analysis section below describes how we used these categories and what we found.

Other production measures calculating

We calculated several ratios or other units of production. We calculated the gas-to-oil ratio (GOR) in standard cubic feet per stock tank barrel of oil (scf/bbl) for each entry and the percentage of gas wasted. Oil and gas are measured in separate units and sold for disparate values, but for much of the analysis we wanted to reduce production rate to a single value, so we created two equivalencies. We used the standard energy equivalence of 6 thousand cubic feet (Mcf) to a "barrel of oil equivalent" (boe), and we used a rule of thumb for value equivalence at 20 Mcf per "barrel of value equivalent" (BOV).

Disposition authority reformatting and binning

The revised rules specified pairs of letters that companies should place in the comments field of the disposition reports to indicate the authority they are relying on, and that authority should be specified for both venting and flaring separately when gas was both vented and flared. Consistent with the format of the forms, the information on authorization authority scraped from the web query was contained in two data fields, both formatted as text. Although the codes were standardized, the overall structure of the comments containing the codes was not. Use of the data required first two normalization steps.

First, we had to algorithmically extract and separate the authorization codes for venting and flaring. The two fields sometimes included labels to indicate whether the authorization code applied to venting or flaring, and sometimes they included extra information. We concatenated the fields, searched for key words, and extracted the text following those key words and assigned the resulting text to separate data fields for venting and for flaring. We tested and verified the algorithm on the full set of unique values in the data set to be sure it worked as intended.

It turned out that in about 10% of entries (each entry a combination of lease and month), operators listed more than one kind of authority for the wasteful disposition of their gas. So, in the second step, we simplified the authorities into a smaller set of unique values. When two compatible authorities were cited, we assigned the more common of the two (i.e., if the operator cited but Exemption and Allowed by Rule, the volume was classified as Allowed by Rule, which is more than 10 times as common. If the operator cited both Authorized by Exception and Exception Pending, then we assigned Authorized by Exception, which is also more than 10 times as common). Still, a small

number of cases contained incompatible codes. For the purpose of understanding the frequency and significance of authorizations, we excluded these cases.

Corporate family standardizing

We observed that many operating companies were listed separately from closely related companies. For example, we observed that companies which had been purchased in total by a larger company retained their own separate identity in the records of the RRC, so we attempted to alias over 7,600 operating licenses into corporate families. We did not bring to bear information on corporate officers or ownership that is available from the Secretary of State, but we did consider the name, mailing address, and phone numbers.

As with the text fields for disposition authority, the addresses of registered operators were widely inconsistent. For example, a suite number could be labelled “Suite” or “Ste” or be left unlabeled. After diligent effort to normalize and parse inconsistently reported address information, we sorted the directory several ways to look manually for separate entities sharing similar names and/or contact information. When companies had enough in common, especially mailing addresses, we deemed them to be related entities.

We merged our resulting analysis into the disposition data sets, showing whether each was a part of a corporate family as well as the name of the company or family. Note that the directory licensees did not include all historical license holders, the aliasing covers only a fraction of the historical disposition data. Also, this aliasing was not used to count transfers of operatorship because the aliasing did not include historical companies.

Reservoir names standardizing

For the sake of estimating gas production by analogy as discussed below, we wanted to identify wells producing from the same reservoirs or geologic plays. Each well is assigned to a field, each with its own rules managed by the RRC. The field name is generally a combination of a surface name and a reservoir name. (Pending wells are yet assigned to a field, and some old fields do not include a reservoir name.) However, the reservoir name can be spelled or abbreviated differently in different fields. The Cotton Valley formation, for example, may be officially identified as “CV” or might be misspelled as “Coton Valley.”

For the sake of choosing analogs, we attempted to normalize all of the variations of reservoir names into a smaller, more consistent set based on our experience and data about how the names were applied. We exported all of the unique reservoir names active in the state as well as some basic info about each, such as the number of wells, the total production of oil and gas, average depth, the primary counties, and the primary regions. Then we sorted to put similar reservoir names together and manually assigned a reservoir alias name based on the official names and related information. The assignment does have limitations, of course, particularly among the relatively small group of reservoirs identified with nothing more than a number.

Analysis and Interpretation of Reported Disposition

Investigating reports of no gas produced

Given that some gas is always produced in association with oil, reporting zero gas production is clearly incorrect. We deemed unreliable every record that reported oil production but no gas production.

Records that did report gas volumes were subject to secondary tests. Given the state of Texas allows operators to estimate volumes and the lax enforcement of reporting at all and given our experience in operations and reporting, we searched for a number of possible patterns that would indicate artificial—and thus likely unreliable—data as described below.

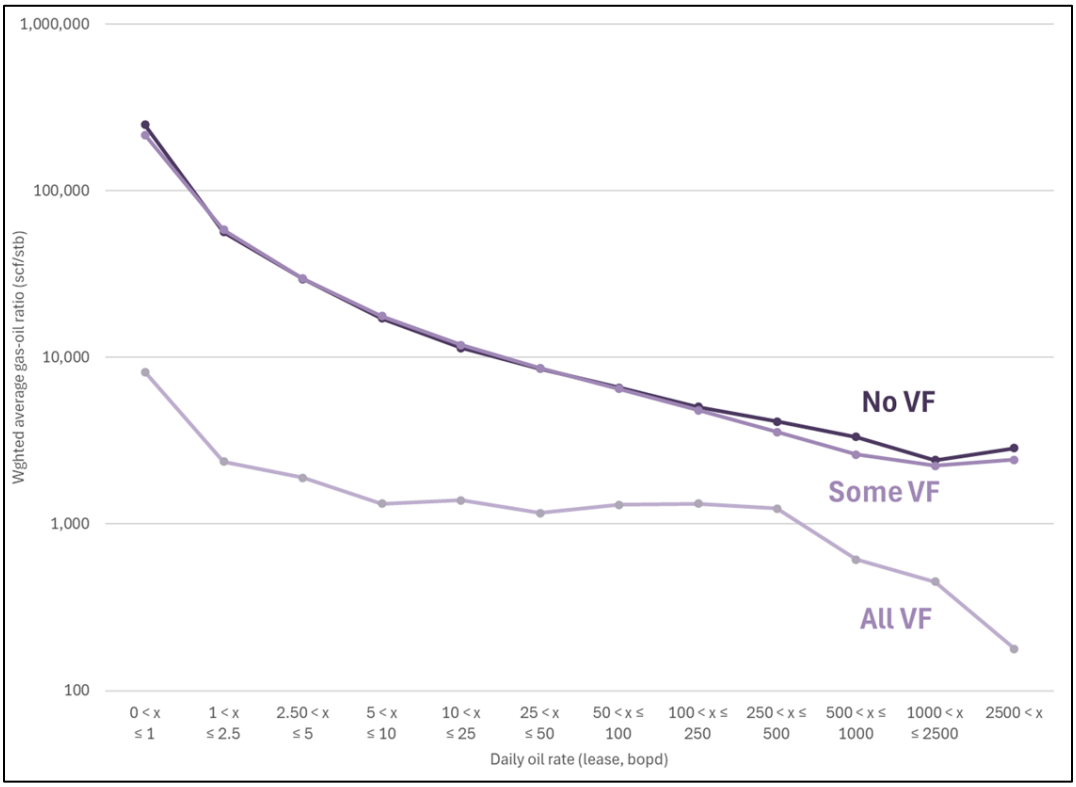
Investigating reports of all gas vented or flared

Rules of reporting allow volumes to be estimated in many cases instead of rigorously measured, and there is no means of accountability or verification for released volumes. Consequently, we expected that released volumes might be underreported, especially when the operator does acknowledge that some oil or gas have been vented or flared.

To test the possibility, we compared the trend of reported GOR with oil production rate at the lease level. First principles suggest that in most cases GOR should increase monotonically as liquid production diminishes. As fluids are removed from the reservoir, oil saturation decreases, and more gas evolves out of the liquid phase. The increased saturation of gas makes it flow much more easily than the liquid (increased relative permeability). In gas reservoirs with condensate, the reduction of reservoir pressure causes liquids to form out of the gas while still inside the reservoir rather than at the surface. Relative permeability effects again mean that less condensate arrives at the surface, and the ratio of gas to liquid increases.

As shown in Exhibit 2 below for statewide data, we graphed the trends of GOR versus oil production rate for different kinds of reporting: no vented or flared, some vented or flared, and all vented or flared. The reasonable trend and consistency between records with no or only limited venting or flaring suggests that these categories are mostly reliable.

Exhibit 2. Statewide relationship between reported GOR and reported daily oil production, separated into groups based on the degree of self-reported venting or flaring (VF).



This graphic includes both oil leases and gas leases (with some oil production). When considering oil leases separately from gas leases, trends of production when some venting or flaring (Some VF) are reported diverge in opposite directions below about 25 barrels of oil per day (bopd) from the trends when no venting or flaring is reported (No VF). Together, the patterns reinforce, as discussed below, that the division between oil leases and gas leases is more regulatory relic than reservoir reality. In most smaller groups examined, we found that the two styles of reporting agreed reasonably well, though the cases with some venting or flaring more often demonstrated lower gas-oil ratios than higher. When this pattern occurs, it suggests that data may be less reliable among cases that acknowledge some releases.

More importantly, the expected pattern of GOR breaks down and diverges in cases for which the operator reports that all gas is vented or flared. Smaller subsets showed consistently and significantly lower gas production when the operator reported all volumes vented or flared. Consequently, we deemed all reported volumes unreliable (and recalculated an inferred rate by analogy) when all gas reportedly vented or flared.[v]

Investigating reports of repeated volumes produced

There still may be unreliable data in cases where some gas is sold or used. Production volume can also be artificial when no volume is reported vented or flared. Sometimes, a company will report exactly the same volume of gas produced each month or for each day in the month, and repetition

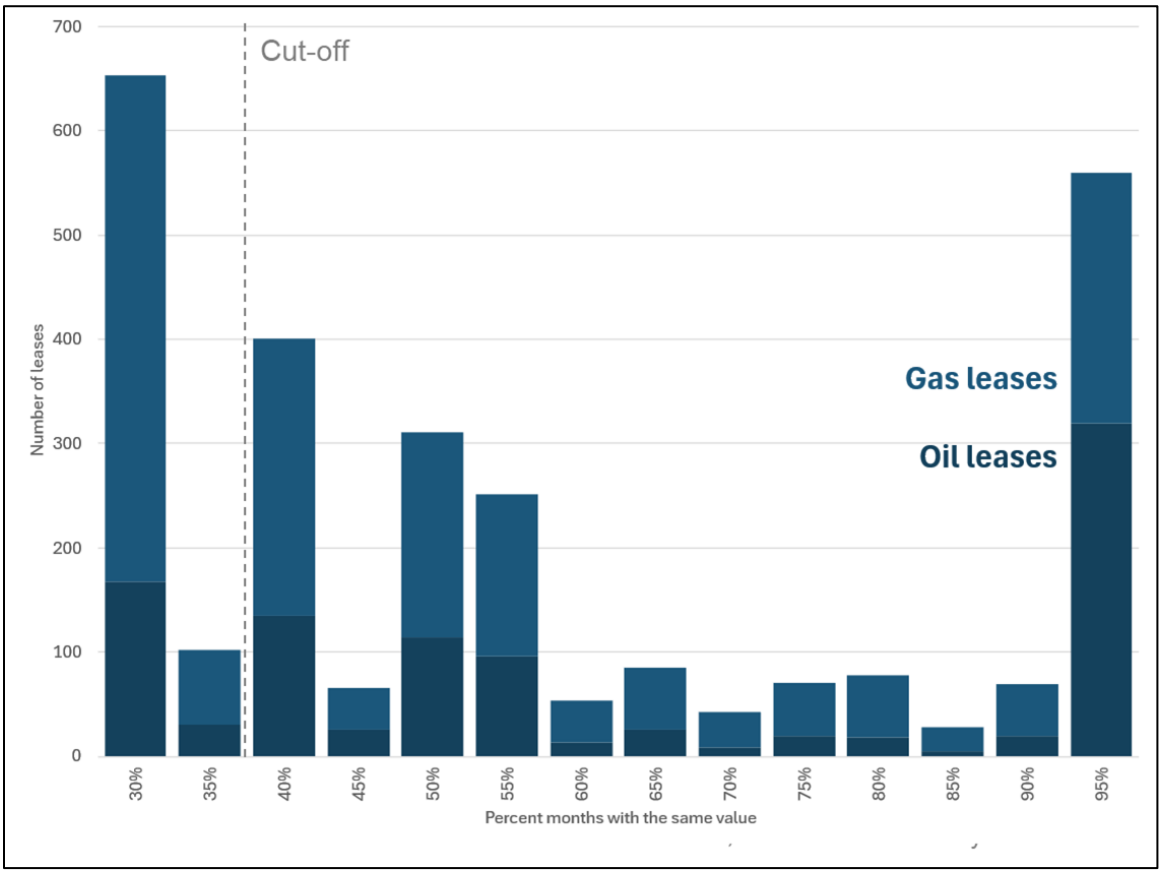
[v] A similar pattern likely exists when all gas is reportedly used on the lease, but we did not examine this subset separately.

of the same figure suggests that the figure is artificial. E.g., an operator may report 1 Mcf every month or may report 28 Mcf in February and 31 Mcf in March.

For all cases, we calculated the number of repeated monthly values and compared it to the number of months with values of gas reported to create a figure for the percentage of months repeated. Specifically, the rates from each lease ID were counted then rank-ordered. A counter stepped through the list from lowest to highest, resetting each time the value changed.[vi] Finally, we found and extracted the rate with the highest counter assigned. If more than one figure was repeated the same number of times, we used the largest of them. Also, leases reporting the same volume of gas every day for 12 operating months would thus report the same total volume in months with 30 days and appear on the graph as having 42% repetition.

We found that it was possible for rates to be repeated a second time by chance and that many of the low-repetition cases contained high flow rates very early in life, so we deemed those more likely reliable (or at least reasonable). Three months of exactly the same volume is very much less likely. Still, we elected to treat any case with three repeated values as reliable, leaving only 0.6% of leases with potentially unreliable data. Among these, Exhibit 3 below shows the percentage of months for which the same value was repeated.

Exhibit 3. Frequency of repeated values of monthly gas production when four or more months shows the same value.



[vi] Note that we did not test separately for months with the same daily rate, only for months with the same total volume. The algorithm could be modified to increment when the monthly volume differs by more than 1/30th.

In the end, we applied two criteria to designate data as unreliable: four or more months of the same value and that accounting for at least 40% of the reported values. However, when this criterion was met, we deemed all of the reported values to be unreliable because we observed that among the high-repetition cases, it was common for less-repeated values to be interspersed with repeated values. So, we elected to categorize as unreliable all volumes from leases with greater than 40% of reported values being the same number.

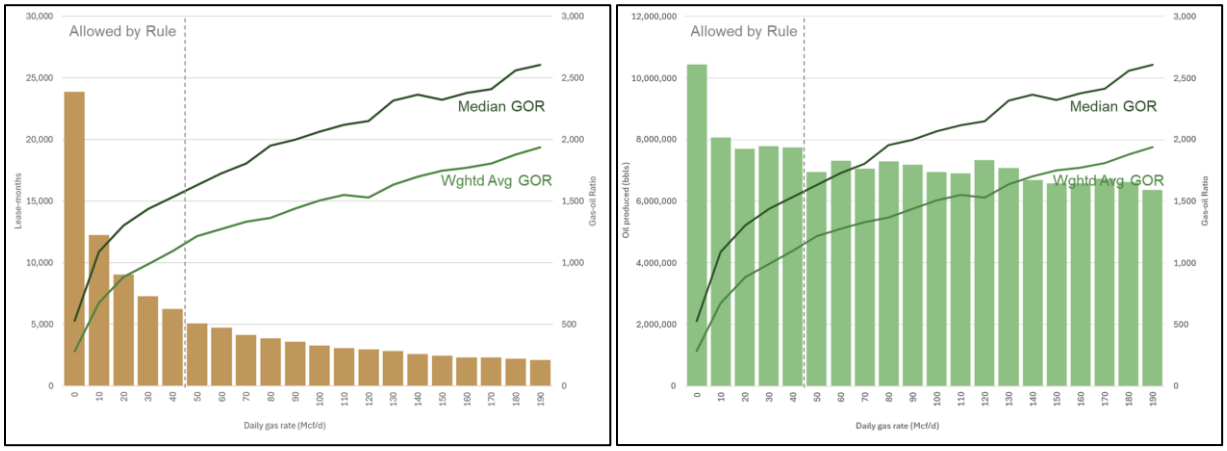
Further tests of reliability were conducted, but in the end no other tests were used to classify records as unreliable. Still, the same tests could prove valuable in other studies.

Investigating reported values in proximity to the regulatory limit

The RRC rules authorize venting up to 50 Mcf per day for each oil lease (regardless of the number of wells) and up to 15 Mcf per day for each gas lease (only one well). So, we examined the frequency of reporting of releases near these thresholds on the theory that production might be artificially estimated below the regulatory allowance.

Exhibit 4 below shows the trends of gas production from oil leases statewide by six measures, but only leases claiming that the release is allowed by rule. In all six graphs, each bin represents a 10 Mcf/d range up to 200 Mcf/d. The blue bars in the first row of graphs show the number of cases in each bin, which is the same in all three graphs. The three graphs, however, show different measures of gas-oil ratio as a line. The first shows the median GOR of each bin, the second shows the average GOR, and the last shows the weighted average GOR. The second row of graphs shows the oil produced in each bin by the height of the bars, again the same in all three graphs. The lines show the same GOR trends as in the first row of graphs.

Exhibit 4. Number of records and associated volume of oil by daily rate of oil production (bars) compared to two measures of GOR.



We do not observe an anomalous decrease/increase in the number of cases or the volume of oil reported above/below the 50 Mcf/d threshold. What is more, the lines on each graph show that the gas-oil ratio also follows a consistent trend without discrepancy around the rule's threshold.

In this statewide data and other large groups, we did not observe irregularities in the patterns of reporting. As we focused on smaller subsets of records to isolate instances, however, the distribution of reported volumes became more erratic.

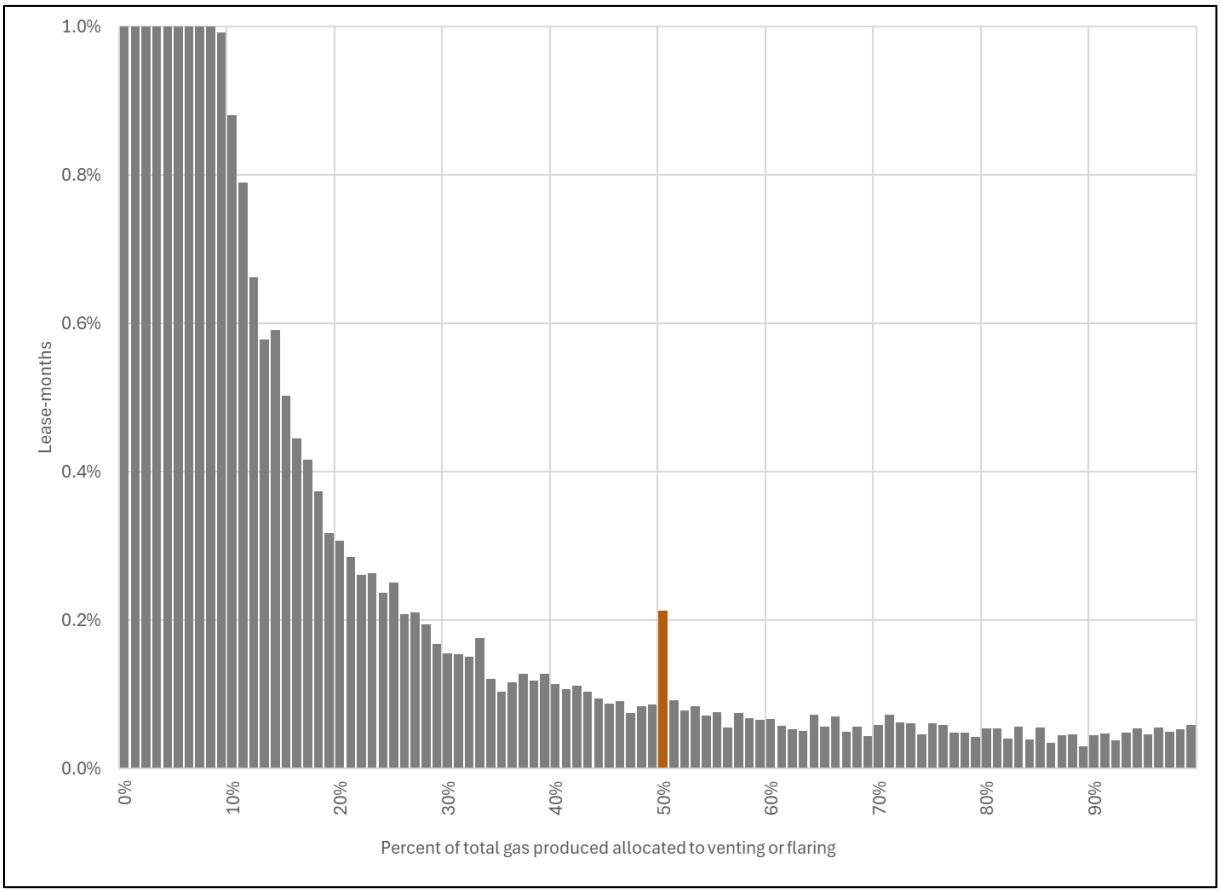
In the end, we could not systematically separate false positives from true positives, i.e., concentrations just below the threshold due to random noise instead of unreliable reporting. Our Benford analysis (below) also demonstrated similar shortages of numbers above the regulatory thresholds for certain companies. However, we did not find a reliable way to identify unreliable figures systematically and algorithmically through the data.

Investigating round figures for percentage of resource wasted

After looking for patterns of artificial volumes produced, we examined some ways that the estimate of volume released may be unreliable. First, because the RRC allows released volumes to be estimated, we hypothesized that releases might be assumed merely as a fraction of overall production without a quantitative foundation. This kind of artificial data turns out to be rare. This follows from the fact that nearly 99% of reported waste is flared and the expectation that flared volumes are more likely to be measured than vented volumes.

We calculated the percentage wasted (vented plus flared) versus total gas disposed of to look for artificial estimates of wasted gas. Note that we did not look for repetitions of the volume released, though either of these might also demonstrate estimates without a quantitative foundation.

Exhibit 5. Histogram of relative frequency of reported percentages of total gas produced reported disposed of by venting or flaring.



For the records with some gas reported to have been wasted, the x-axis of Exhibit 5 shows the percentage of the total shown to be wasted binned to single percentage points. The y-axis shows the

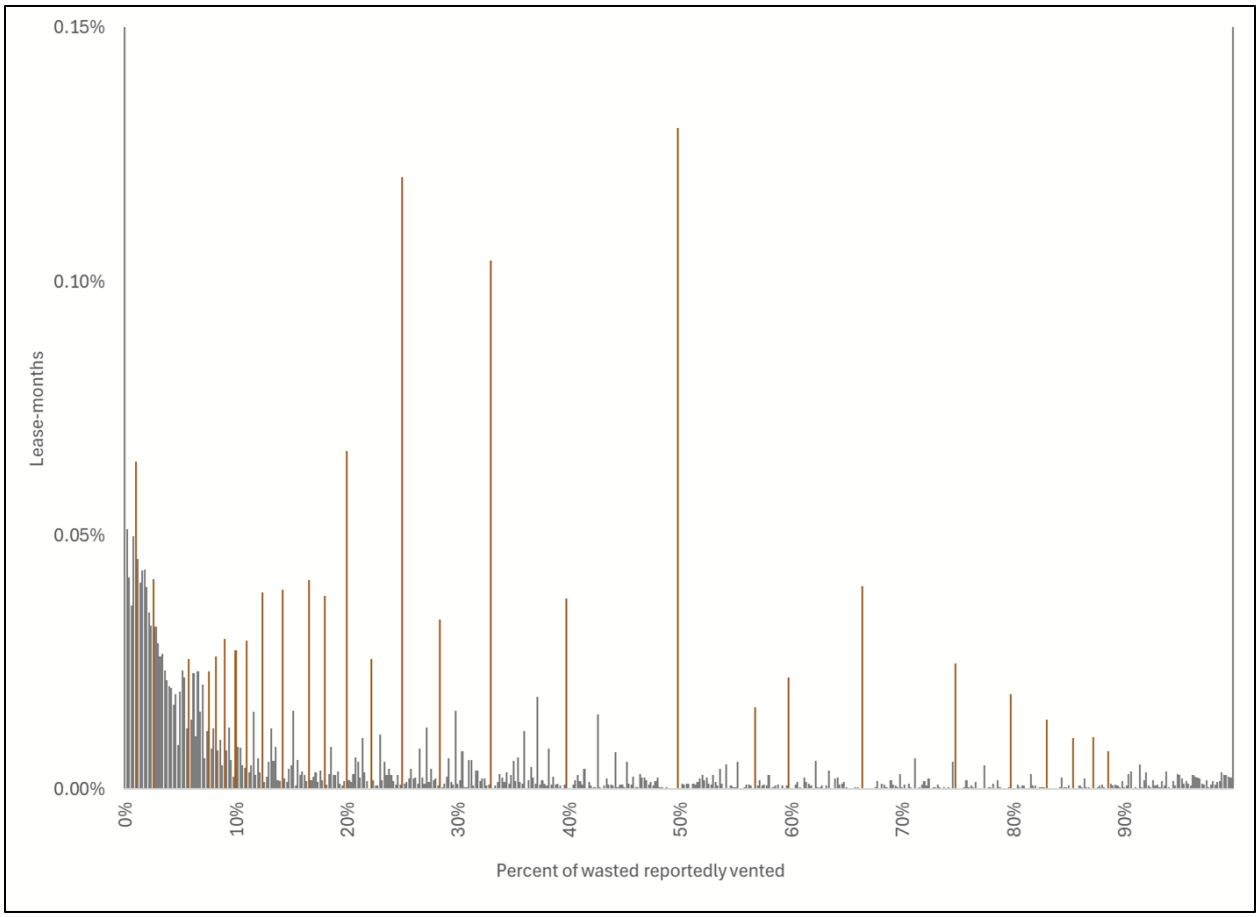
number of cases in each bin, but the scale has been truncated to focus on the lower range of counts. When some gas is wasted, the large majority of cases show less than 10% wasted, and the trend of fractions is mostly smooth and reasonable. There is a clear anomaly at 50% and a less clear anomaly at one-third. They represent few cases overall, and those are concentrated around just a couple of operating companies, so we did not flag these for correction.

Investigating round figures for split between vented and flared gas

The next level of detail looked within the same cases that did report some venting or flaring but examined the split between reported venting and reported flaring volumes. As above, we did not look for repetition of absolute volumes of vented or flared. Given the paucity of venting among gas leases, we examined only oil leases. Of course, with so few cases reporting venting at all, the population of samples was small, but the data does demonstrate anomalies.

Exhibit 6 below shows how operators split wasted gas between vented and flared in records with at least some venting. The x-axis shows the volume reportedly vented as a percentage of the total volume reported wasted (i.e., vented + flared), binned to show the fraction precisely. And the height of each bar indicates the proportion of records (lease months) in the bin.

Exhibit 6. Distribution of relative frequency of the percentage of wasted gas released by venting among records with only some gas wasted.

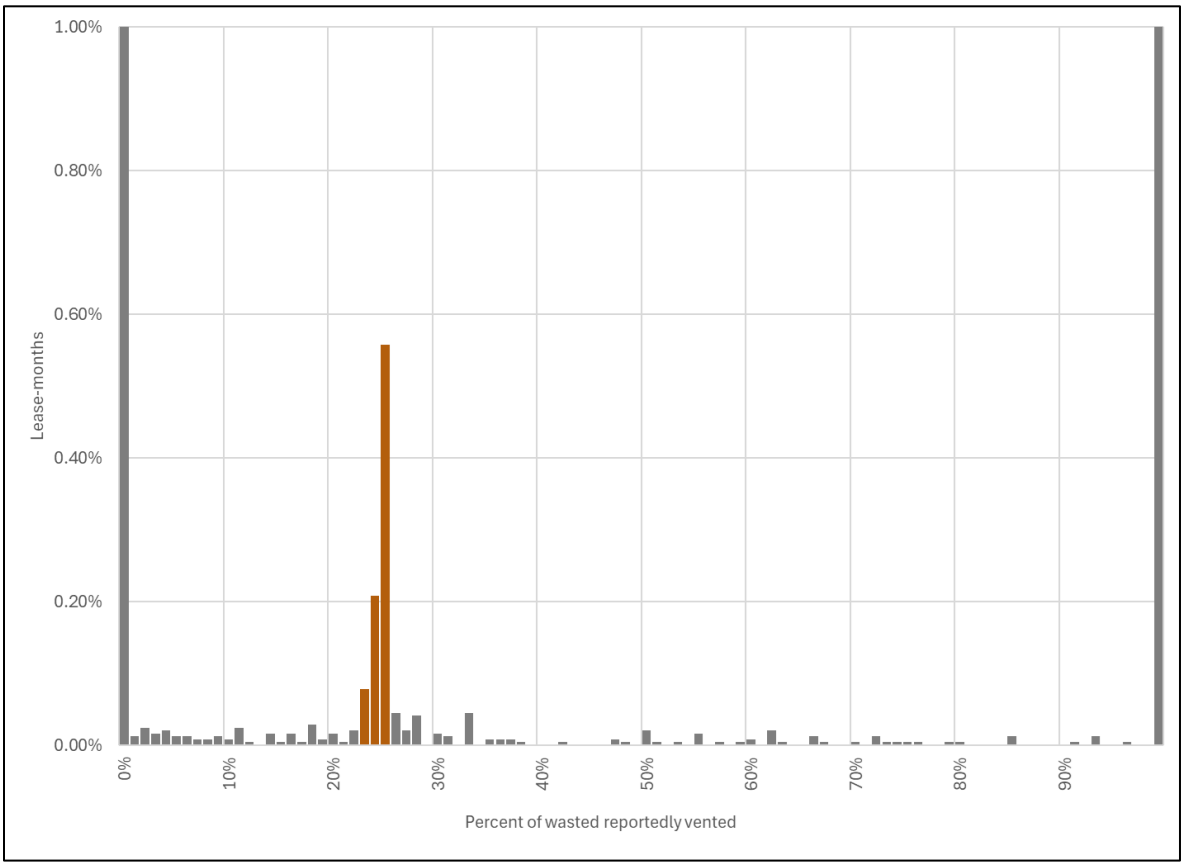


The large majority of records show either 100% vented or 0% flared. Not many split between the two, but there is a strong pattern of round numbers in between the limits. Most of the discretization

appears to be driven by the precision of the self-reported data, specifically the division of low integers. Unlike some other artificial data patterns, a number of companies reported in this pattern, including many large companies. For example, a major company which recently acquired and now operates old wells in the Spraberry Field of the Permian basin is responsible for largest part of the anomalously round figures.

We also looked at the split between types of waste when all of the gas is reportedly wasted. Exhibit 7 below also shows the percentage of gas vented for the group of records in which all gas is reportedly wasted. As when only some gas is vented for flared, the vast majority of cases report either all vented or all flared. Only about four companies attempt to split between vented and flared when all gas is wasted. One company accounts for the spike of anomalous data observed at 33% vented.

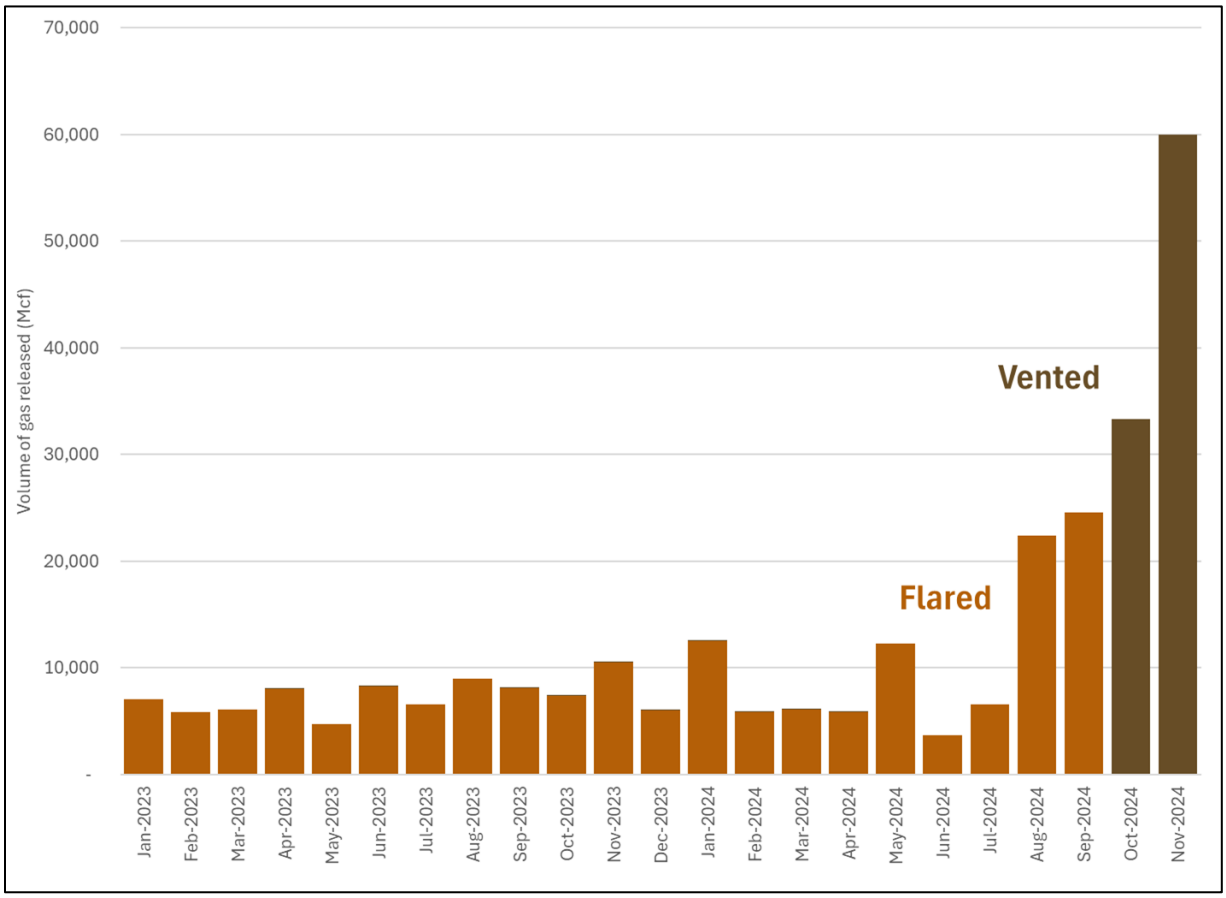
Exhibit 7. Distribution of relative frequency of the percentage of wasted gas released by venting among records with all gas reportedly wasted.



Investigating statistically significant step changes in reported releases

In a few cases, we observed sudden and significant changes in the volumes reported to have been released. For example, Exhibit 8 below shows the reported volumes reported flared (lighter color) and vented (darker color) by PDC Permian ("PDC") for its wells in the Phantom field of the Delaware basin monthly from January 2023 to November 2024. Chevron Corporation purchased PDC in May 2023 and closed the acquisition in August 2023. Notwithstanding this change of corporate ownership and control, the wells were still operated under the name of PDC Energy and would not appear on a standard search for wells operated by Chevron.

Exhibit 8. Monthly volumes of gas reportedly released from wells operated by PDC Permian in the Phantom Field of the Permian basin.



The reported volumes demonstrate two step changes. From July to August, a year after the transfer to Chevron, the volume reportedly flared increased by about three times. The lack of transition suggests a step-change in reporting rather than a change in operations around the first of the month. Then, from September to October, the volumes increased again, and they converted from reportedly all flared to reportedly all vented. In this case, the disposition changes entirely at the start of October, but the volumes released increase in both October and November, suggesting the possibility of a change in field operations.

Step changes like this could reflect a change in field operations or merely a change in reporting practice. We did not attempt to find or interpret — much less to modify — these kinds of anomalies. Research into a particular operator or field merits a review for this kind of step change in reporting.

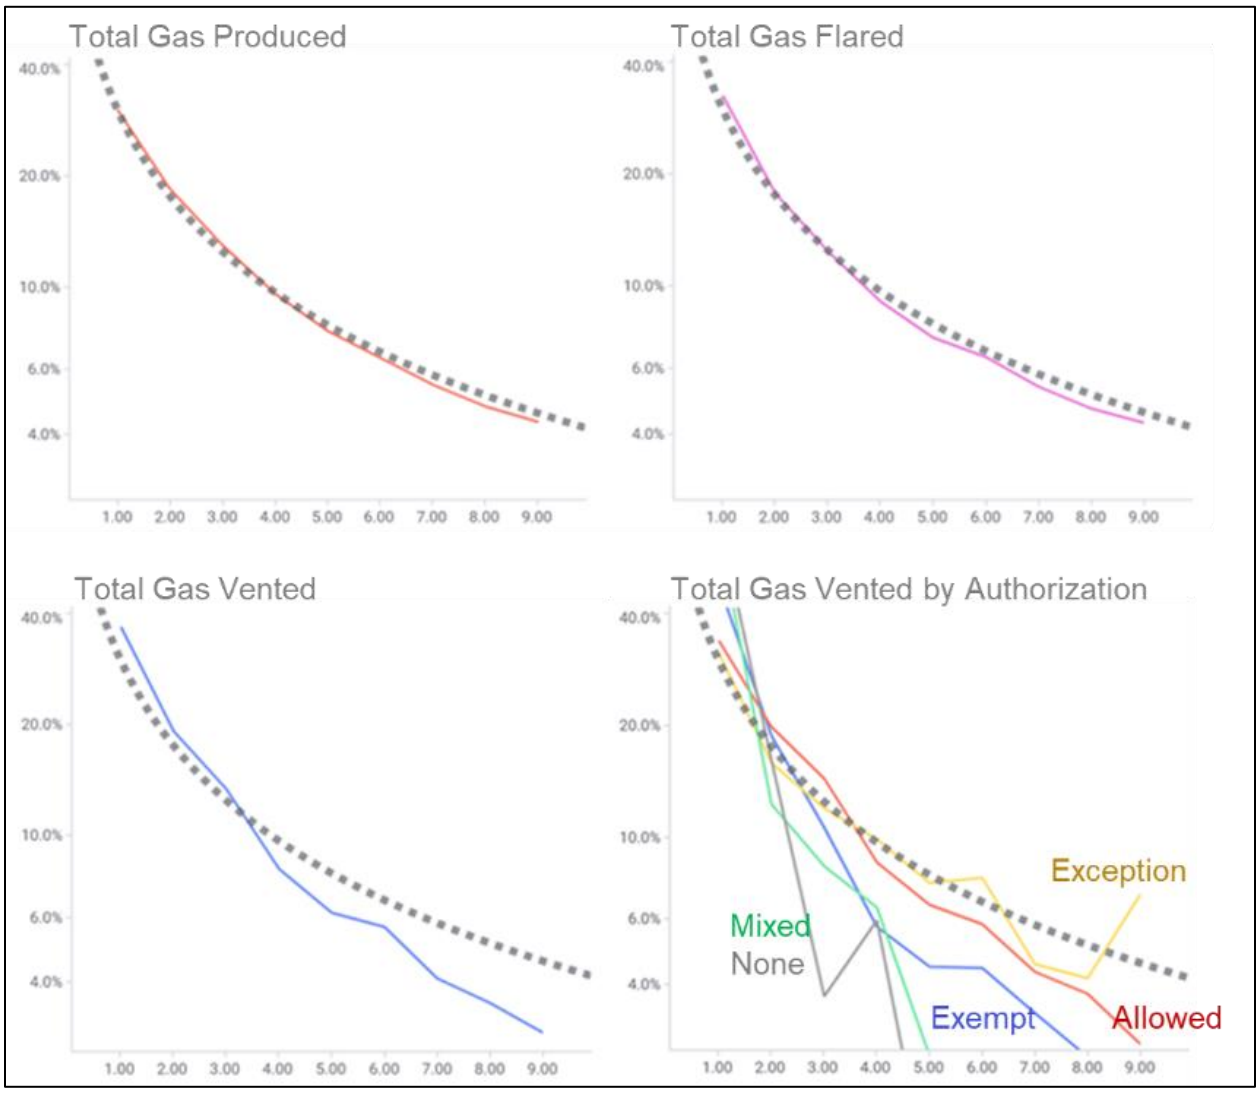
Benford's Law analysis of reported gas disposition

Benford's Law analysis compares oil and gas self-reported data against the probability distribution of leading digits. If the numbers are random and spread over several orders of magnitude, then the proportion of numbers as leading digits should follow a trend established by Benford's Law. Non-random number patterns appear on a graph as deviations from theory. Hence, financial auditors and others use the analysis to look for artificial or fraudulent data. The analysis looks for systematic errors and informs the analyst where to dig more deeply, but it does not by itself identify which

numbers are artificial. In this case, Benford analysis reinforced the same conclusions demonstrated above.

Exhibit 9 below demonstrates one such pattern. The thick gray dashed lines represented Benford's Law for the predicted distribution of random numbers, and y-axis is logarithmic so that the distance between the law and actual represents a proportional difference than an absolute difference.

Exhibit 9. Testing reported gas disposition using Benford's Law.



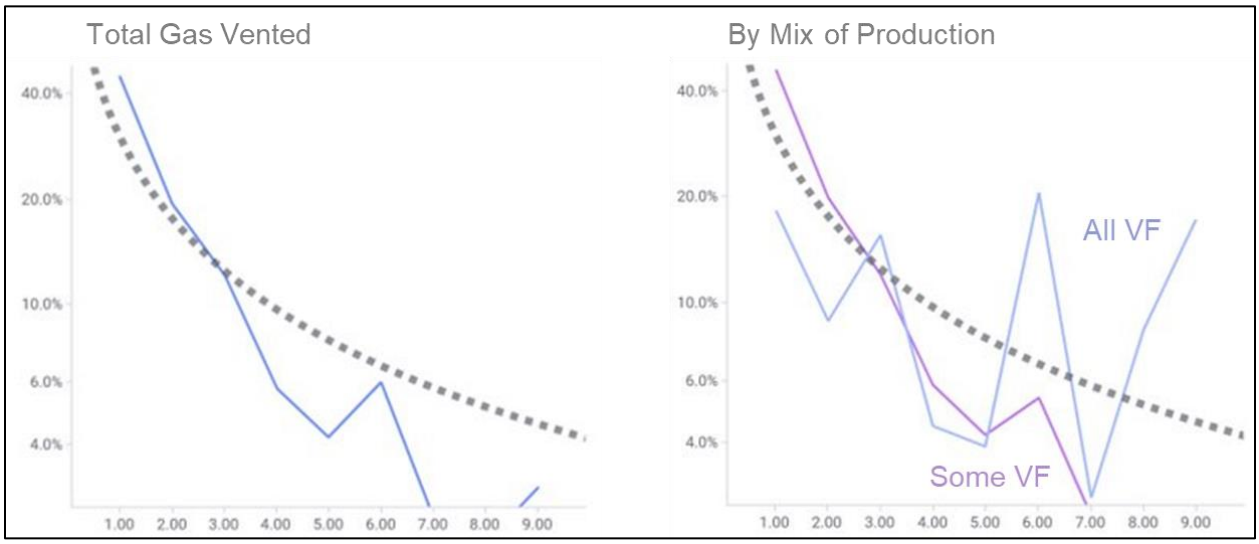
The top left graph shows the distribution of gas reportedly produced. It is by far the largest dataset, and it follows the theory well. Oil production and disposition similarly followed the law except that leading digits of 1 were over-represented because it is not allowed to report fractions of a barrel. The distribution truncates at 1, and a significant number of cases report either 0 or 1 each month. In the upper right, the leading digits of gas reportedly flared also comport with theory.

The bottom two graphs, however, show how the reported volume of gas vented deviates. This is a much smaller population of numbers, but it does show that the figures seem more reliable when the operator must apply for permission to vent ("Exception" shown in yellow), than when it is

automatically allowed ("Allowed" in red and "Exempt" in blue). In the few cases when the operator reports contradictory authorizations or does not report as required, the figures deviate even more.

Many companies operate small portfolios or similar wells which can limit the utility of Benford's Law analysis, but we did find it useful in a couple of cases, such as shown in Exhibit 10 below. These graphs show the pattern of reported venting by one of the largest public operating companies in the state. The left side shows a significant deviation from the law using more than 4,000 records. The graph at right breaks the same data into two groups: when all gas is reported vented or flared ("All VF") and when only some is reportedly released ("Some VF"). This variation shows the lower-quality reporting when all gas is released, and it also shows three anomalous concentrations of monthly volumes around 30, 60, and 90 Mcf, representing 1, 2 and 3 Mcf/d.

Exhibit 10. Testing of reported gas disposition from a single operator using Benford's Law.



Benford analysis of individual operators also sometimes suggested patterns of systematic reporting below the regulatory thresholds for allowed releases, for example, with a dramatic shortage of gas volumes released beginning with the digit 5.

Extension to Actual Gas Volumes Released

It became necessary next to estimate both the actual volume of gas produced, and its disposition. To make these estimates, we interpreted the disposition data deemed reliable (without bringing in additional data) to estimate the released volumes.

Oil is the more valuable commodity and the easier to handle, so we deemed measurements of oil volumes to be reliable in all records. We estimated the unreported gas releases by assuming the same gas-to-oil ratio (GOR) reported for analogous leases with records deemed reliable.

We had to make three choices: how to calculate GOR, how to choose among analogs, and how many analogs were enough (or too many). We evaluated sensitivity in all three dimensions. In the end, all calculations use the same method of calculating gas-to-oil ratio (GOR, scf/bbl), and all use the same bins of oil production rates to choose analogs. Variation of these inputs was not justified.

Most of the methodology focuses on prioritizing the choice of analogs and the minimum number of cases for the application of each, and we report the results of the sensitivity analysis. We also ran two large-scale alternative scenarios; one based only on production from wells designated by the RRC as oil wells and one with all wells that produced liquid hydrocarbons.

Following is an explanation of considerations and conceptual design. Farther below is a more concrete step-by-step explanation of the data and calculations made in the spreadsheet. But first one note on the premise of the calculations.

Determining quality of "reliable" data

It should be noted that the data deemed reliable — and used to estimate releases — probably underreports production and releases. The incidence of unreliable data and the incidence of questionable data both increase at low oil production rates, so it follows that similar problems exist more widely but in ways we could not identify. For example, lower production rates correlate to more releases exempt from estimation and reporting. Lower rate wells are generally older wells more likely to suffer from degraded and leaky equipment, i.e., to suffer greater fugitive releases. And lower producing wells tend to require more repairs and thus more releases from blowdowns and purges of equipment necessary before repairs. Unquantified issues with "reliable" data likely bias our data to the low side.

Determining GOR calculation basis

Because it is a ratio, the distribution of GORs tends to be wide and skewed, specifically to the high side. For a particular population, the arithmetic average of individual GORs tends to be greater than the total GOR (total gas divided by total oil). What is more, GOR can vary from month to month with normal volatility of production and reporting.

Because we are trying to estimate total gas, we did not average the GORs observed among the individual analogs or individual months. Instead, we used the total GOR of the group of analogs calculated over an entire year of reporting. That is, we assumed that the GOR of a subject lease was the same as the total gas produced by the choice of analogs divided by the total oil produced. And those totals were totals represented not individual months, but the entire year of data examined.

Determining analog selection approach

To choose among analogous leases, we identified, examined, and chose among three types of criteria: type of lease (i.e., oil or gas), geologic similarity, and liquid production rate.

Type of lease: The RRC categorizes producing wells as either oil wells or gas wells, and it regulates the two categories differently. For example, production from oil wells is reported in groups while gas production is reported for each well separately. Similarly, the RRC authorizes gas releases by rule for both designations, though at different rates.

At the highest level, we chose two groups of analogs: oil leases (as designated by the RRC) and all leases (oil and gas) that reported non-zero liquid production. The first uses a smaller set of analogs with more similar histories. The second represents a realistic alternative high side because the physical behavior of the reservoirs, especially late in productive life, is more continuous than binary. Wells produce a mixture of hydrocarbons regardless. The RRC differentiates between oil leases and gas leases (primarily) based on reported gas-oil ratio; any lease with GOR over 100,000 scf/bbl must

be classified as a gas well. We have already established, however, that the reported gas volume is not reliable, so the production could be as analogous to gas leases as to oil leases.

Geologic/Reservoir similarity: Different reservoirs contain different fluids and deplete under various kinds of producing mechanisms. Both of these strongly affect GOR, both are strongly correlated to geology. Geology, in turn, is correlated to location. We tested a wide range of criteria for reservoir similarity including county, district, statewide; RRC definition of field and reservoir, field, aliased reservoir, geologic structural feature; the combination of geologic structural feature and reservoir, the combination of geologic basin and aliased reservoir.

We preferred to use closer analogs when enough data was available, and we preferred more intrinsically geologic (if somewhat interpretative) criteria. We grouped and ordered the choices of analogy from preferred to less preferred, then we ran sensitivity cases on each. In the final calculation, we ignored location-only criteria and chose minimum thresholds for each of the preferred analogs. The step-by-step outline below provides more explanation.

Liquid production rate per lease: We created a third, independent dimension for the choice of analogs based on oil production rate of the lease. As depletion lowers the pressure in a hydrocarbon reservoir, the production rate decreases while the ratio of hydrocarbon gas produced to hydrocarbon liquid increases. So, we generally expect gas-to-oil ratio to increase monotonically with lower production rates, and we deemed this a first-order variable.

Reported volumes range over five orders of magnitude, from less than 1 to more than 1,000 barrels per day, with a rate reported down to the individual barrel in each month. This distribution is not uniform or even symmetrical, and it has no inherent discretization. It is a nearly continuous distribution with the bulk clustered at low rates and a tail skewed right.

It must be noted that the disposition and production data are reported by RRC Lease ID and that leases categorized as oil-producing can include any number of wells. (Gas leases each include only a single well). Ideally, we would bin production on the basis of single-well rates, and we did find records from the RRC to show how many wells on each lease were still unplugged. However, the public data does not show directly whether unplugged wells are actively producing or standing idle, and it is known that a large proportion of wells are standing idle. We chose to select analogs on the basis of daily production from the lease instead of daily production from the unplugged (active plus idle) wells on the lease. This choice requires fewer assumptions or interpretations, and we do not have a reason to expect it to be biased in one direction or another.

As the distribution of lease rates is roughly logarithmic, we binned production rates using geometric bins adjusted to fit within a base-10 logarithmic scale. After some testing that the results behaved as expected, break points were finalized at 1, 2.5, 5, 10, 25, 50, 100, 250, 500, 1,000, and 2,500 bopd as shown in Exhibit 10 above.

Determining minimum number of analogs

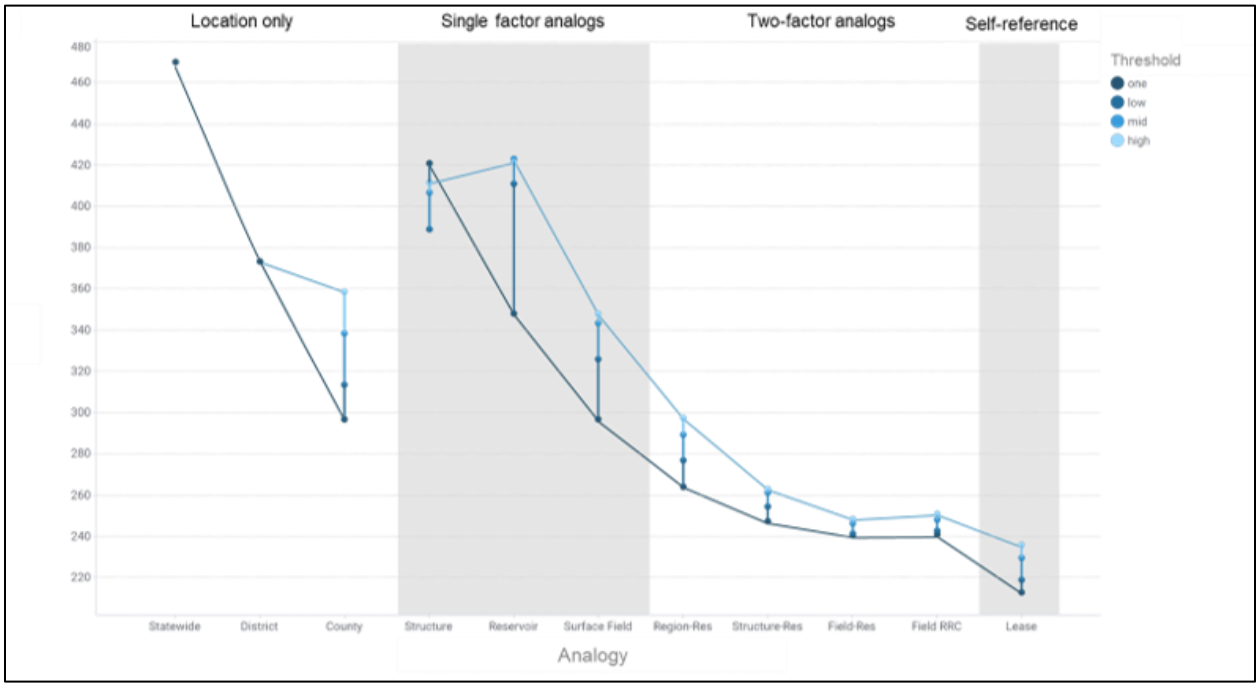
Of course, the quality of the analogy should initially improve as the number of analogs increases. As described below, we tested a range of thresholds for the minimum number of cases for each choice of analog. In the end, we chose the same minimum threshold for each level of analogy.

The long-high side tail of GOR distribution described above suggests that the distribution may also be skewed with respect to the weighted average GOR which we use in the calculation. In other words, a greater share of the probability mass may lie below the true average than above it. As a result, a small data sample is more likely to underestimate the average GOR than to overestimate it.

On the other hand, it is possible that a larger set of analogs also includes more heterogeneity and thus less similarity. The skew of the GOR population could, therefore, result in a weighted average that is too high as it oversamples less similar cases of gassier production.

In practice, we observed that higher minimum thresholds almost always result in higher estimates of release and that wider analog groups consistently result in higher estimates of gas released. We have not been able to parse these dynamics. We also observe that the Spraberry (Trend Area) and a few other fields in the state are quite large and diverse. Future work could refine the secondary input of minimum threshold, and future work could consider also breaking down such fields into smaller groups so that the analog data does not become too broad.

Exhibit 11. Range of inferred production based on scope of analogs (x-axis) and minimum threshold for use of analog (data series).



Determining whether gas is vented or flared

Unreported gas was obviously not sold as it would have been measured and easily reported. Similarly, volumes of reinjected gas deserve to be measured for reservoir management purposes. Both the lack of measurement and the low oil rates associated with the majority of unreliable values (as discussed more below) suggest that the unmeasured gas was more likely vented.

It is possible that some of the unreported gas was used on location. Self-reports show that overall lease use of gas runs just a few percentage points, though self-reporting of lease use at low rates seems unreliable. As a practical operational matter, use of gas on site becomes more difficult at low oil production rates (which accounts for most of the inferred volumes) because flow is intermittent.

It is more likely that at least some of the inferred gas produced was flared instead of vented. On the one hand, volumes flared are easily measured in most cases, and flaring is more needful and more practical than venting at higher gas flow rates. On the other hand, third-party satellite observations document approximately twice as much flaring as was reported to the regulator. The question

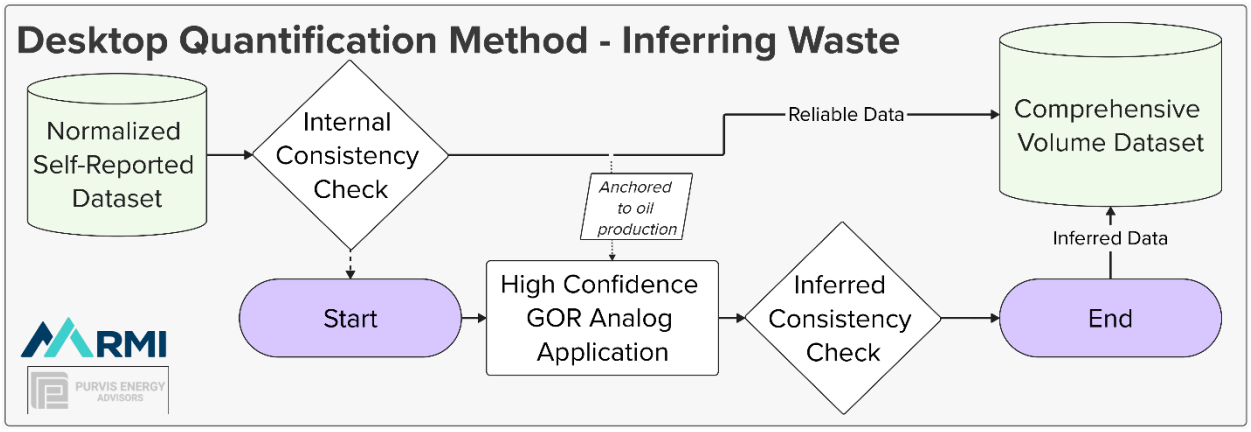
becomes how much of that observed-but-unreported flaring comes from leases we deemed reliable and how much from leases for which we inferred a corrected volume, and we have not examined this question on a lease-by-lease basis.

It is possible that the excess flaring came from leases which our algorithms deemed reliable. That is, the reported figures may look reasonable when they are in fact systematically low. It is also possible that some of the unreliable data for which we inferred the produced volume was actually flared. This possibility is more likely at higher production rates. For reference, the difference between self-reported flaring and observed flaring is approximately equal to our inference of unreported volumes produced on leases reporting more than roughly 35 bopd. The majority of the inferred production, however, comes from leases reporting less than 15 bopd.

On balance, we believe that the majority of the unreported volumes were disposed by venting. For the sake of validation as described below, we consider alternatives.

Step-by-step for Inference of Actual Volumes

Exhibit 12. Process flow of inferring waste within overall methodology.



The calculations of corrected values were made in Microsoft Excel, then later ported back to TIBCO Spotfire for quality control and validation (see below). Inputs to the calculations — namely, info about leases and about groups of analogs — were aggregated for each lease over the 12 months being studied. That is, we make the calculations on an annual basis instead of a monthly basis.

We exported information about all of the leases as two tables, one with header information and one with production data. Because header data fields with identifying information are text and because some, like the name of the operating company, can change, we aggregated the information to show the most common value among the 12 months of records for each lease in the base data. The header information included columns for all of the data fields of analog groups.

Columns of production data were grouped into reliability sub-categories: to be used as analog, to be replaced because all gas was vented or flared, to be replaced because no gas was reported, and to be replaced because of repeated values. For each of these groups, we exported the sum of oil production, the sum of gas production, and the number of months with non-zero production. In MS Excel these separate exports were combined into a single table, we calculated their daily oil

production rate bin considering the volume produced and the months on, and we calculated the amount of oil produced for which a new estimate of gas was needed.

For each category of analogs (such as county, field or reservoir), we export one table of summary data. The monthly data was filtered to include only data deemed reliable as an analog and aggregated over the 12 months being studied. Each line represented one group of analogs, and columns contained sums of oil produced, gas produced, and the number of producing months. However, sums of production information were separated into different columns based on the bin of the oil rate. E.g., if a lease reported 45 barrels of oil in a month, then its 45 barrels were summed into the column of oil production for the rate bin of 1 to 2.5 bopd. The associated gas would be summed into the column for gas produced from oil production in the bin of 1 to 2.5 bopd, and the well count associated with the same bin would be increased by one. The exported totals of oil and gas volumes were combined in the spreadsheet to calculate the GOR for each analog (row) and bin (column).

Next, we matched the analog GORs to the leases by looking up the lease header information in the summary tables of analogs and returning the GOR and the number of cases included in that average GOR. This analog information was arranged from the most specific analog (on the left) to the most general analog (on the right).

At the very beginning of the sequence of analog GORs, we calculated the GOR of production from the lease which was deemed reliable. That is, if a lease had a mix of reliable and unreliable records during the study window, we calculated the GOR of the reliable production and the percentage of the months which were deemed reliable. This became the first choice of analogy—other production in the same year from the same lease—but it applied only to the cases in which the operator reported that all gas was released and reported the volume released.

At this point, the calculation contains actual and analog GOR for every lease and every type of analog, creating a complete set of data for quality control in the next step.

To calculate a single answer, we set up a cell for the minimum number of cases necessary in each analog category, and a formula chose the most precise analog available with the prescribed minimum number of analogs[vii]. Finally, the chosen analog GOR was multiplied by the volume of oil needing new estimates, and the reported production volumes were subtracted from the new estimate of gas production to arrive at the volume of gas released but not reported.

It should be noted that this algorithm does allow for negative corrections. That is, if the methodology predicts less gas was released than was reported, then the correction is negative and offsets upward corrections in other leases. Negative corrections were, however, only about 1% of corrections, demonstrating the systematic nature of underreporting.

In order to determine the minimum threshold for use of an analog, we used the Data Table function of MS Excel to test a range of minimum values. We isolated each analog by setting the minimum threshold for preceding analogs above the actual maximum set size and setting subsequent thresholds to 1, then we calculated statewide results at multiple thresholds. Then we selected a single minimum threshold for each analogy. Besides helping to choose the deterministic thresholds, these calculations also serve as a way to measure the quality of each analog as a whole. Also, as noted above, we did not apply a maximum threshold, so a few of the more precise analog sets include larger number of presumably more diverse analogs.

[vii] Sample data and a formula template is included with this methodology.

The calculation of inferred gas production was performed on two subsets of the statewide data: leases categorized as oil leases and leases that produced non-zero volumes of oil. The analog inference was not performed on all wells, so the tabular outputs of the calculation had to be merged with the results previously deemed reliable to create a complete, corrected view of production in the state. Attached is a sample of our calculations in MS Excel and of the final recombined data set.

Validation of methodology, internal controls

During execution of the defined methodology, we validated internal consistency across datasets and investigated any discrepancies prior to proceeding to subsequent steps. This validation included quality control (QC) within individual methodological steps, quality control across steps, and holdout tests for inferred data. In the end, a kind of holdout (or out-of-sample test or blind test) was conducted to examine the results of the methodology against higher-confidence volumes. Specifically, we compared what the methodology would predict for each of the cases used as input to the average analog value.

QC within individual methodological steps

The nature of the separate data tables was examined before loading and use. For example, we examined ranges of values, lists of unique values in a field, repetition of values we assumed to be unique, and lists of all records sorted by value. As described above, we sometimes found small sets of unexpected data which needed to be adjusted or removed before integrating with other data.

QC across individual methodological steps

When the same data was available from more than one source, we compared the values from the separate sources. If the overlapping data did not align, we could not expect that the remainder of the merging table would apply to the base table. While merging data in SQL for the first part of the study, we compared totals, made spot-checks of individual leases, and compared the unique values in the data fields. For the calculation of inferred production in MS Excel, we counted and totaled columns, reproduced calculations by hand, searched for and resolved error messages, and analyzed the implied gas-oil ratios at a summary level.

QC via reproduction of higher-quality data

To validate the calculation of inferred volumes, we exported intermediate and final calculations of inferred production to Spotfire where we visualized graphs and tables to compare against self-reported figures deemed reliable.

The procedure consisted of the following steps after modeling inferred production in MS Excel.

1. Ported reported and modeled production data as well as intermediate steps to Spotfire
2. Filtered leases with reliable reported production data.
3. Cross-plotted estimated production volumes against reported volumes
4. Made pivot tables to compare aggregate modeled production against aggregate reported production.

Our calculation of inferred gas production relies on the weighted average GOR of analogs, but for the back-testing we compared the starting point and ending points, namely the volume of gas produced to the volume of gas predicted to have been produced. The objective of the model and the significance of the results are volume of gas; the ratio is merely an intermediate step.

We found that, as expected, applying the single average of analogs across all cases reproduces the average of the analogies overall but leaves wider variations in individual leases and smaller subsets. Individual leases show considerable variance from reported, but over- and under-estimated trends tend to offset within larger groups and push the average error toward zero. As a natural corollary, we did find that the smaller-scale analogs created narrower ranges of errors than large-scale groups, so we still prefer the more precise analogs.

The blind testing confirmed the large-scale accuracy, but it also confirmed the variance of individual lease estimates. The results can be totaled as precisely as a single operator with a single lease, but they should thus be considered more representative when explaining larger groups. The estimate will be less precise for smaller groups.

Validation of methodology, external controls

To evaluate the results of the methodology, we compared model outputs against independent, external estimates of emissions and loss ratios in the state. Some comparisons were straightforward, but others were less direct. Within the range of uncertainty, our work compares favorably.

Comparisons to prior estimates need to consider these dimensions of the measurement.

- Date – Our analysis pertains to November 2023 to October 2024. We compare against studies of emissions from, mostly, 2022 and later.
- Area – Some studies look only at a basin (or the extent of a shale play within a basin) while some look at statewide totals.
- Type of waste – Studies may examine only methane emissions (mainly venting), only flaring, or total wasted as venting plus flaring.
- Scope of oil and gas activities – Natural gas flows from wells through gathering systems to processing plants through pipelines to refineries, and all these components are considered part of the "oil and gas industry." Plus, fugitive natural gas can leak from inactive and plugged wells.
- Units of measure – Our work measures emissions in volume of gas while others report the mass of methane and/or the loss ratio of methane.

The most uncertain element of comparison comes from different or potentially different scopes of oil and gas activities included. Other studies sometimes report emissions from well sites separately, but more often they include gathering and some subsequent steps. Our work calculates emissions based on active oil production and implies that it is released within the control of the upstream operating company. These same companies commonly, but not always, also control gathering infrastructure within the field.

As discussed above and again below, there remains a question of the disposition of the volumes we infer to have been produced. For the comparisons below, we try to adapt our inferred volumes as

we can to match the types of emissions in other studies, and we present alternative disposition types when we can.

It is not clear whether other investigators assigned to gas gathering should be included in the comparison to our work. It is clear, however, that our work does not include emissions from gas processing plants, transmission pipelines, or refineries, and it does not include emissions from idle or plugged wells.

Units of measure of that vented gas also introduce a potential variance. To compare measurements of methane mass to natural gas volumes requires, most importantly, an assumption about the methane content of the natural gas. We have assumed methane constitutes 74% of the produced natural gas, a methane density of 0.657 kilograms per cubic meter, and thus that 13.8 kg translates to one thousand cubic feet (Mcf). The methane concentration across the state is interpreted from a recent study of gas compositions in key US basins.

Comparison of flaring to top-down satellite observation

VIIRS NightFire Flaring and World Bank Global Flaring Partnership Satellite produce observed flare data. Statewide totals account for nearly twice as much volume flared as self-reported by oil companies. It is possible but unclear whether the observed-but-unreported flares account for part of the unreported gas production which we assumed was vented.

NASA and NOAA deploy the Visible Infrared Imaging Radiometer Suite (VIIRS) instrumentation on several satellites[4], and the VIIRS NightFire for Flaring ("VNF")[5] project at the Colorado School of Mines' Payne Institute for Public Policy interprets the size and intensity of infrared frequency emitted by the burning of natural gas to interpret the volume of gas combusted using a well-validated correlation.

The VNF project found evidence of the burning of 6.2 billion cubic meters (bcm) of gas during calendar year 2024 across the state of Texas, equivalent to 219 Bcf. Similarly but less precisely, the World Bank Global Gas Flaring Tracker Report, July 2025 used the same data plus additional sources to conclude that flaring in the Permian and Eagle Ford basins totaled 237 Bcf in 2023 and 233 Bcf in 2024.[viii] By comparison, self-reported volumes flared totaled 118 Bcf during the detailed study period of November 2023 to October 2024. The empirical evidence shows about twice as much incinerated gas as self-reported to regulators.

It is well-documented that not all flaring is self-reported to regulators, and this dynamic likely dominates the discrepancy, but two other dynamics could explain part of the difference between self-reported and observed volumes of flaring.

Observations of burned gas should, in theory, be less than the volumes measured on their way to the flare because not all gas routed through flaring equipment is, in fact, flared. Not all flares are lit, and combustion of a lit flare is not complete. When the gas is not ignited, it may be measured and reported as being flared even though it is not actually combusted. The best research available suggests that about 10% of "flared" volumes are not actually combusted (Plant et al. 2022) and thus not observed by VNF. This dynamic would make observed combustion (VNF) estimates lower than self-reported volumes, the opposite direction of the primary error.

[viii] These figures do include part of New Mexico, but New Mexico has stricter rules on flaring natural gas than Texas. And they do not include flaring in other active parts of Texas such as the Haynesville shale.

It may be noted that 90% combustion on 118 Bcf of self-reported flaring implies 12 Bcf vented instead of flared while the 90% combustion estimate applied to observed combustion would imply 23 Bcf intended to be flared but instead vented. Of course, the discrepancy is much larger and in the opposite direction.

The discrepancy might or might not overlap with our estimate of unreported production volume. It is possible that VNF observes the flaring of unreported volumes which we have estimated, and this dynamic would shift inferred volumes from the vented category into the flared category. It is also possible that VNF observes flared volumes associated with leases which we identified as reliable. That is, the observed-but-unreported flare volume might come from the leases we examined or might not.

Comparison of flaring plus venting to bottom-up hybrid measurement statewide

Our estimate of total volume wasted (flared plus vented) compares favorably to bottom-up hybrid modeling published by Climate TRACE.

OCI+ calculates and compares emissions intensities across oil and gas assets in Texas, elsewhere in the United States, and internationally, using a methodology independent of the analysis presented here. The methodology relies on information about the nature of operations, data-based relationships observed in the real world to produce results. Where available, the OCI+ uses open-access data in consistent, comparable, and verifiable formats, though due to the limited publicly available information, the modeling efforts mostly rely on proprietary data purchased from data providers. Alongside standardized datasets, RMI integrates in-house geological and petroleum system expertise to improve the accuracy of subsurface inputs. Some of the open-access data used is reported production. OCI+ estimates are used to quantify methane waste and other greenhouse gas emissions as inputs to Climate TRACE, and its results are shown in the next section.

Using this analysis as input for Texas operations in calendar year 2024, Climate TRACE reports approximately 6.8 Mt of methane wasted per year as venting or flaring. Assuming methane constitutes 74% of the produced natural gas, then 13.8 kg translated to one thousand cubic feet (Mcf), and Climate TRACE implies 493 Bcf of natural gas flared or vented. Self-reporting totals only 120 Bcf released, but our work suggests 375 Bcf were flared or vented in the base analysis and 551 Bcf in the high scenario.

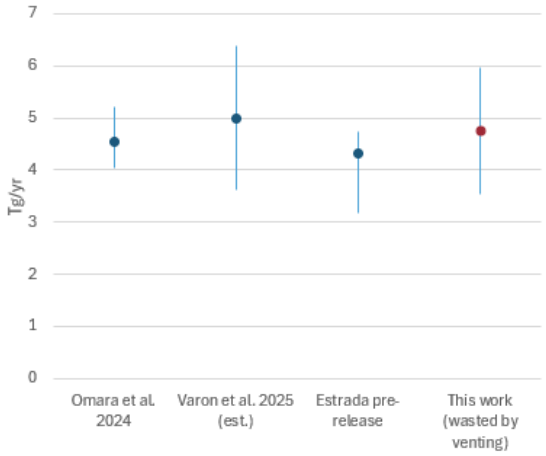
Our figures fall within the range of statewide estimates of the mass of methane emitted as presented by and interpreted from other recent studies relying on empirical observations of emissions, though ours could be on the high side of the range of other studies is gas gathering losses are treated differently.

Among the statewide studies presented here and the basin studies presented below, some studies categorize empirical measurements of methane emissions from different kinds of facilities or equipment, treating the population of measurements as a sample and extending the observed distribution of emissions to all of the facilities in the state. Other studies reverse engineer the volume of methane released based on empirical measurements of excess methane in the atmosphere as measured by fixed sensors on towers, aerial flights, and satellite measurements. Of

course, this inversion process requires other inputs, sometimes including emissions inventories based on self-reported data.

At 13.8 kg/Mcf, our estimate of venting translates to 3.5 to 6.0 teragrams (Tg) of methane vented from upstream operations in the 12-month period. Exhibit 13 below compares this range against the date, scope, and uncertainty of other statewide estimates of methane emissions from the oil and gas production. Overall, our analysis is consistent by leaving room for other kinds of emissions such as fugitive leaks and pipeline blowdown events. The figure represents our best understanding of the data available, but the sources and figures merit some discussion.

Exhibit 13. Comparison of measurements and calculations of methane released each year by oil and gas operations in the state of Texas.



Study			high	low	Tg/yr
Omara et al. 2024	extrapolated sample	(CY 2021)	5.20	4.04	**4.53**
Varon et al. 2025 (est.)	satellite inversion	(CY 2023)	6.4	3.6	**5.0**
Estrada pre-release	satellite inversion	(CY 2024)	4.7	3.2	**4.3**
This work (wasted by venting)		(2023-24)	6.0	3.5	**4.8**

“Estrada pre-release" comes from pre-lease text and data visualization publicly available from the author for an update to the analysis previously published by Varon et al. 2025. Both studies use satellite data to improve the analysis of the US Environmental Protection Agency’s Greenhouse Gas Inventory, which is known to underestimate releases from oil and gas production. Also, the work uses a modified inversion algorithm which reduces estimates of emissions in the Permian and becomes less accurate there in order to become more accurate in other parts of the country with lower emissions.

Varon et al. 2025 reports totals for only a few of the main basins in the state, a total of about 4.3 Tg/yr with a high side confidence interval of about 5.3 Tg/yr. By visual comparison to Estrada pre-release over the entire state, we estimate that the previous work by Varon would imply roughly 5.0 Tg/yr in the base case. Similarly, “EDF unpublished” comes from personal communications with the methane emissions team at Environmental Defense Fund based on their analysis protocols, including MethaneAIR and MethaneSAT measurements.

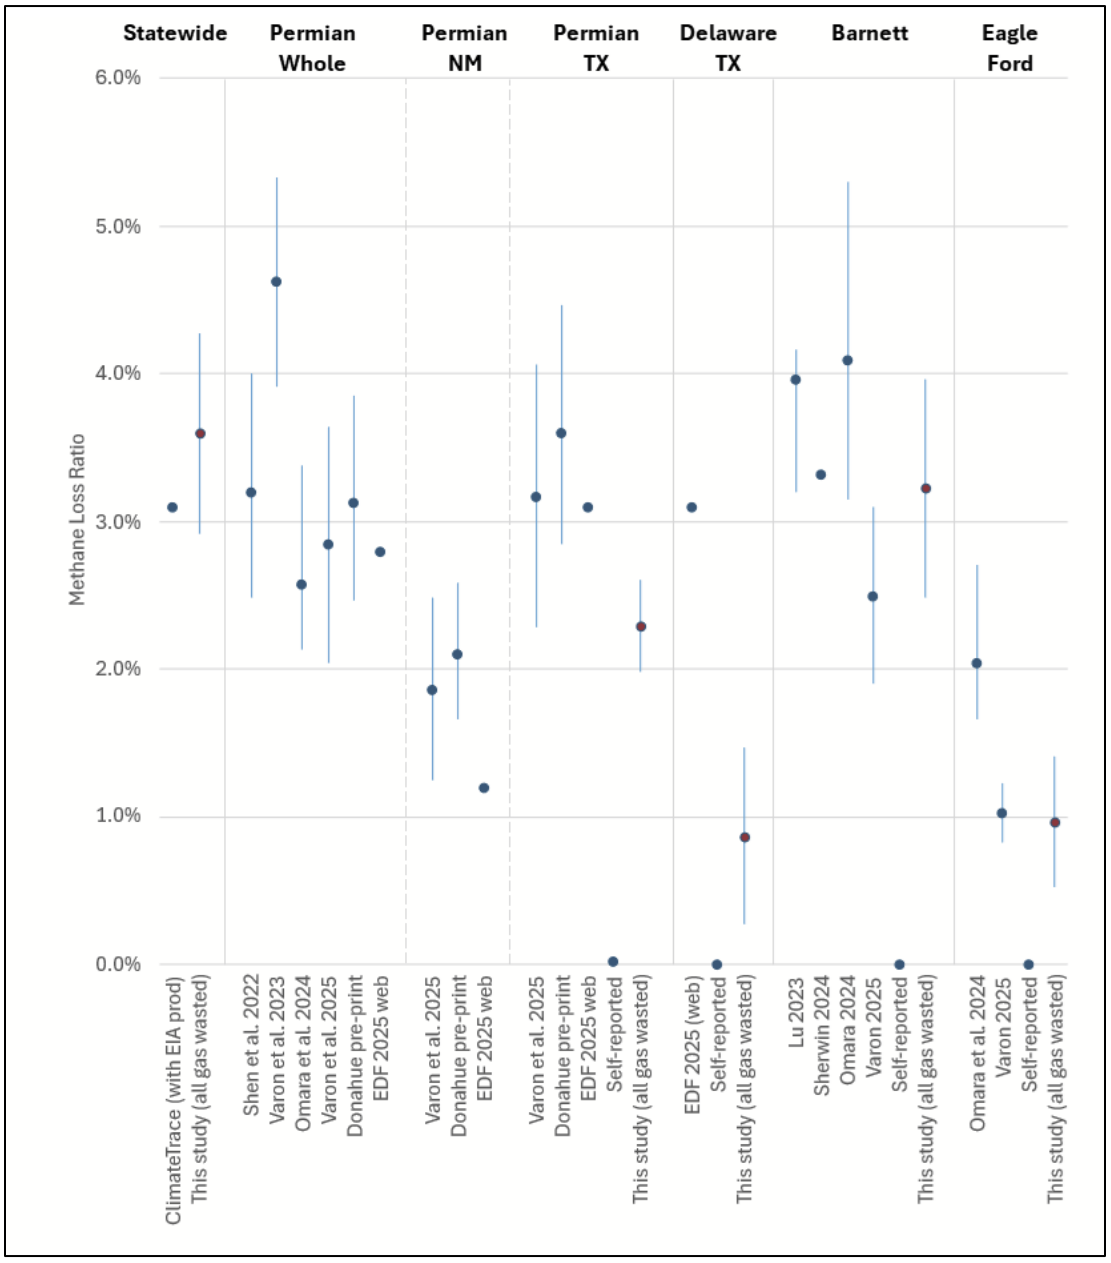
Our work fits within the range of previous results, but it should be noted that if it is more appropriate to remove emissions from gas gathering, then the other studies would decrease by 15% (explicit in Omara et al. 2024) and 30% (by analogy to other studies discussed below), and our results

would be on the high side or above other studies. On the other hand, if our inferred volumes do not include observed-but-unreported flared volumes, then our work implies venting mostly higher than these analogs.

Comparison of venting to empirical and semi-empirical loss ratios by basin

Much of the literature reports emissions by producing basin (or part of a basin) and often as a loss ratio (which is less sensitive to variations in defining the boundaries of the basin). Though the question of emissions from gas gathering persists, our work also compares favorably to these studies. Exhibit 14 below demonstrates the comparisons.

Exhibit 14 Comparison of measurements and calculations of methane loss ratio in major basins of Texas.



The figure does involve some interpretation or adaptation. The studies do not all use exactly the same basin boundaries, do not all report loss ratio as such, and do not all use the same assumption of methane content. It may be noted also that the calculation of loss ratio in our work uses the same denominator as previous studies, namely reported gas production and including our estimate of unreported gas production.

After normalizing the results as best we can, we found that our results comport well with previous empirical and semi-empirical results. Statewide, our figure for total gas vented and flared comports well with the ClimateTrace model of releases normalized to statewide production as reported by the US Energy Information Agency.

On the Texas side of the Permian basin, our figures come in similar to but somewhat lower than empirical measurements even when we assume that all over inferred volume is vented. In the areas of the Barnett Shale and the Eagle Ford Shale, our figures are also similar or perhaps low. To the extent that some of the inferred releases are actually flared, our loss ratio would be lower and allow more for the possibility of other sources. It may also be noted that in all of these areas where third-party studies found 1.0 to 4.1% methane lost, self-reported figures ranged from 0.003 to 0.02%.

Previous studies, and thus this table, focus preferentially on areas of marquis development. Combined with the fact that higher production correlates to lower emissions, studies have often focused on areas with proportionally *lower emissions*. Outside of the Permian Basin and the Eagle Ford play, our study shows the average waste ratio to be 2.9% to 5.6%.

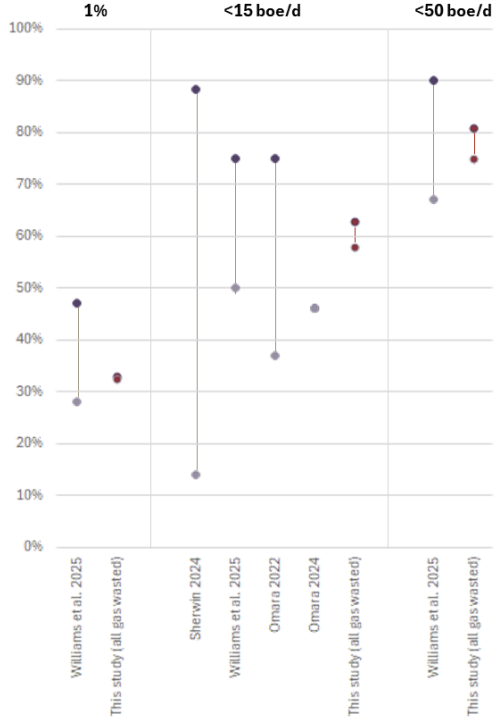
Exhibit 15 Tabular version of comparison of measurements and calculations of methane loss ratio in major basins of Texas.

Basin	Study			high	low	Methane Loss Ratio
Statewide	ClimateTrace (with EIA prod	operational model	(CY 2024)			**3.10%**
	This study (all gas wasted)	corr. reporting	(2023-24)	4.28%	2.91%	**3.60%**
Permian Whole	Shen et al. 2022	satellite	(CY 2019)	4.00%	2.49%	**3.20%**
	Varon et al. 2023	satellite inversion	(CY 2019)	5.33%	3.91%	**4.62%**
	Omara et al. 2024	extrap. sample	(CY 2021)	3.38%	2.13%	**2.58%**
	Varon et al. 2025	satellite inversion	(CY 2023)	3.64%	2.04%	**2.84%**
	Donahue pre-print	aerial plus extrap.	(CY 2024)	3.85%	2.47%	**3.13%**
	EDF 2025 web	combination	(2024-45)			**2.8%**
Permian NM	Varon et al. 2025	satellite inversion	(CY 2023)	2.49%	1.24%	**1.87%**
	Donahue pre-print	aerial plus extrap.	(CY 2024)	2.58%	1.66%	**2.10%**
	EDF 2025 web	combination	(2024-45)			**1.2%**
Permian TX	Varon et al. 2025	satellite inversion	(CY 2023)	4.06%	2.28%	**3.17%**
	Donahue pre-print	aerial plus extrap.	(CY 2024)	4.46%	2.84%	**3.60%**
	EDF 2025 web	combination	(2024-45)			**3.1%**
	Self-reported		(2023-24)			**0.02%**
	This study (all gas wasted)	corr. reporting	(2023-24)	2.61%	1.98%	**2.30%**
Delaware TX	EDF 2025 (web)	satellite inversion	(2024-45)			**3.10%**
	Self-reported		(2023-24)			**0.003%**
	This study (all gas wasted)	corr. reporting	(2023-24)	1.47%	0.27%	**0.87%**
Barnett	Lu 2023	satellite inversion	(CY 2019)	4.16%	3.20%	**3.96%**
(Ft. Worth Basin)	Sherwin 2024	aerial direct plus extrap	(CY 2021)			**3.32%**
	Omara 2024	extrapolated sample	(CY 2021)	5.30%	3.15%	**4.09%**
	Varon 2025	satellite inversion	(CY 2023)	3.10%	1.90%	**2.50%**
	Self-reported		(2023-24)			**0.003%**
	This study (all gas wasted)	corr. reporting	(2023-24)	2.49%	3.96%	**3.23%**
Eagle Ford	Omara et al. 2024	extrap. sample	(CY 2021)	2.70%	1.66%	**2.04%**
	Varon 2025	satellite inversion	(CY 2023)	1.23%	0.83%	**1.03%**
	Self-reported		(2023-24)			**0.003%**
	This study (all gas wasted)	corr. reporting	(2023-24)	0.53%	1.41%	**0.97%**

Comparison of venting as a function of production rate

Previous studies have demonstrated by statistical sampling and extrapolation that a large portion or a majority of methane emissions come from low-producing well sites, and our work shows the same pattern. Exhibit 16 below summarized previous studies in comparison to our results.

Exhibit 16 Comparison of estimates of the contributions of methane emissions from low-producing wells.



Production Contribution	Study			high	low
1%	Williams et al. 2025	extrap. sample	(CY 2021)	47%	28%
	This study (all gas wasted)		(2023-24)	33%	32%
5%	Sherwin 2024	aerial direct plus extrap	(CY 2021)	88.2%	14.1%
<15 boe/d	Williams et al. 2025	extrap. sample	(CY 2021)	75%	50%
	Omara 2022	extrapolated sample	(CY 2019)	75%	37%
	Omara 2024	extrapolated sample	(CY 2021)	46%	46%
	This study (all gas wasted)		(2023-24)	63%	58%
10%	Williams et al. 2025	extrap. sample	(CY 2021)	90%	67%
< 50 boe/d	This study (all gas wasted)		(2023-24)	81%	75%

Williams et al. 2025 reported the contributions of the wells delivering the bottom 1%, 5%, and 10% of production nationally. The 5% corresponds to production from "wellsites" producing less than 15 boe/d as defined in previous literature and historical practice as "low producing" wells. Williams and Omara et al. 2022 before him reported the range of uncertainty in the national average of the contribution of these "marginal" wells. Sherwin et al. 2024 reported the range of point estimates from six separate producing basins. Authors have observed that higher proportions prevail in low-producing basins, i.e., basins with fewer young and high-producing wells with low emissions.

For this primary measure of the contribution of low producers, Exhibit 16 shows our distribution of waste gas statewide.[ix] Our two scenarios of higher and lower total emissions do not show much difference in the distribution of sources. And our work falls within the range of estimates for all previous studies that supplied ranges. If we assume that some of our waste was flared instead of vented at higher rates, then the proportion of venting would increase.

The exhibit also compares our statewide average against the national average reported by Williams et al. 2025 at lower and higher production levels. Again, our calculation for the state of Texas comports with the range of national averages previously calculated.

[ix] It may be noted, however, that our work uses production by lease (as reported and recorded in the state) and not individual well (as implied by the previous work).

3. Framework limitations and future development

Limitations

- Ours and other ideas on misreporting should be tested in Texas as well as other jurisdictions. The kinds of misreporting may be different or more pronounced in other jurisdictions, and there may be other kinds of misreporting.
- Information might be collected from completion forms to try to test and/or correct releases made during initial completion.
- Some analog groups are quite large. The county could be combined with a geologic similarity (like field or reservoir) to narrow the scope of analogs for very large groups.

Future development opportunities

- The same calculation could be reproduced for previous periods to reconstruct a history of unreported releases since 1993 in Texas.
- The trends of GOR with oil production rate could be examined in groups of percentage of gas sold. The volume sold should be the most reliable disposition while others, like gas used on location, may not be. Examining this method would isolate cases where most or all of the gas is used in a manner less likely to be measured accurately.
- To the extent that significant reported volumes are arbitrarily split among dispositions, a methodology could be developed to identify and calculate the volume in need of correction. It seems at present that the correction could be made in aggregate without assigning to specific leases or that the correction could be made to those operators demonstrating the pattern.
- Analysis could be made of step changes to reporting of venting and flaring volumes (such as the start of reporting no gas production) to characterize when increases happen, such as what stage of production, what kinds of companies, what years, and what natural gas prices.
- The calculations of inferred gas production could be run in SQL or Python instead of MS Excel for greater integration and simplicity or to be run over historical years.
- Some kind of estimate could be made of the proportion of unreported releases which are, in fact, flared prior to release. For example, the locations and magnitude of observed flares could be compared to reported flaring and inferred venting.
- In other areas with fewer stacked reservoirs (thus less variability among wells drilled physically close to each other), it might be sufficient to choose analogs within a search radius or within a county without interpreting reservoir or geologic structure.
- Oil rates could use an interpretation of the number of active wells to bin by rate per well instead of rate per lease.
- Oil rate production bins could be doubled or halved to test the effect on the range of errors in the blind test.

- The aliasing of current operating companies could be expanded to include corporate information from the Secretary of State.
- Operating company aliasing could be extended to companies no longer active in order to observe corporate trends over time.
- Estimates of releases could be compared against self-reported releases delivered to the federal Environmental Protection Agency, to investors, or to other authorities.
- Self-reported volumes of flaring could be compared against independent satellite measurements.
- In the future, it may be possible to improve the methodology by interpreting RRC data on, e.g., annual tests of individual well production rates to estimate the number of actively producing wells on a lease. Such an interpretation is already made by some data vendors, although its reliability remains unclear.

4. Acronyms and abbreviations

ACRONYM, ABBREVIATION	MEANING	DESCRIPTION
API	Application programming interface	The RRC does not provide an application programming interface (API) or bulk download mechanism for some of the data analyzed
ARCGIS	Arc Geographic Information System.	Software tool used to locate the leases and execute spatial intersections with structural features.
ASCII	American Standard Code for Information Interchange	The data format used in the extraction process of disposition fields and writing the results to a flat file.
BEG	Bureau of Economic Geology	The geologic survey for the state of Texas .
BCM	Billion cubic meters	The VNF project found evidence of the burning of 6.2 billion cubic meters (bcm) of gas during calendar year 2024.
BCF	Billion cubic feet	Self-reported volumes flared total 118 Bcf during the detailed study period of November 2023 to October 2024.
BOE **BOEPD**	Barrel of oil equivalent (or Barrel of energy), Barrels of oil equivalent per day	A unit to combine oil and gas on the basis of energy content, specifically 6 Mcf per BOE.
BO, BOPD	Barrels of oil, Barrels of oil per day	Unit of oil volume or daily oil production rate.
GOR	Gas-to-oil ratio	We calculated the gas-to-oil ratio (GOR) in standard cubic feet per stock tank barrel of oil (scf/bbl).

MCF	Thousand cubic feet	Standard measure of the volume of gas produced, measured at standard temperature and pressure.
MNAR	Missing not at random	Missing data mechanism where missingness is related to the unobserved value itself.
MS EXCEL	Microsoft Excel	The calculations of corrected values were made in Microsoft Excel.
MT	Mega tons	Common unit for the mass of methane released. Equal to 10^6 (one million) metric tons and to 1 terragram.
P-5	Organization Report	Mandatory filing required by the Railroad Commission of Texas (RRC) for entities operating within the jurisdiction of the RRC Oil & Gas Division.
SCF/BBL	Standard cubic feet per stock tank barrel of oil	Unit used to express gas-to-oil ratio (GOR).
RRC	Railroad Commission of Texas	Texas' regulator of upstream oil and gas operations and other aspects of the energy system.
SQL DATABASE	Structured Query Language database	Language used to load, merge, and manage datasets in a relational database.
TG	Terrragram	Common unit for the mass of methane released. Equal to 10^{12} (one trillion) grams and to 1 megaton
VF	Vented or flared	Combination of volumes both vented and/or flared, total gas wasted.
VNF	VIIRS NightFire for Flaring	Project that interprets satellite infrared data to estimate volumes of gas combusted.
VIIRS	Visible Infrared Imaging Radiometer Suite	Satellite instruments deployed by NASA and NOAA used to detect flaring activity.

5. Key Formulas and Conversions

This section formalizes the calculations described in the methodology. All formulas correspond directly to procedures explained in the text and are presented here for clarity and reproducibility.

Gas-to-Oil Ratio (GOR)

The gas-to-oil ratio (GOR) is calculated in standard cubic feet per stock tank barrel of oil (scf/bbl):

$$GOR_{Reported} = \frac{Gas_{Reported}(scf)}{Oil_{Reported}\ (bbl)}$$

Where:

- Gas is measured in standard cubic feet.
- Oil is measured in stock tank barrels.

Percentage of Gas Wasted

The percentage of gas waste (vented plus flared) relative to total gas disposed is calculated as:

$$\boldsymbol{Wasted\ Gas\ Percent} = \left(\frac{\boldsymbol{Gas_{Vented}} + \boldsymbol{Gas_{Flared}}}{\boldsymbol{Total\ Gas\ Disposed}}\right) * \mathbf{100}\%$$

Inferred Gas Production

For leases with unreliable or missing gas data, inferred gas production is calculated using the selected analog GOR:

$$\boldsymbol{Gas_{Inferred}} = \frac{\boldsymbol{GOR_{Analog}}}{\boldsymbol{Oil_{Reported}}}$$

The volume of gas released but not reported is calculated as:

$$\boldsymbol{Gas\ Volume\ Correction} = \boldsymbol{Gas_{Inferred}} - \boldsymbol{Gas_{Reported}}$$

Endnotes

[1] "Permian Basin GIS Projects and Databases," Bureau of Economic Geology, The University of Texas at Austin, accessed January 1, 2026, https://www.beg.utexas.edu/resprog/permianbasin/gis.htm.
[2] William E. Galloway et al., *Atlas of Major Texas Oil Reservoirs*, Bureau of Economic Geology, The University of Texas at Austin, 1983, https://store.beg.utexas.edu/publications/atlases-of-major-oil-and-gas-reservoirs/at0002.
[3] "Texas County Boundaries Detailed," Texas Department of Transportation, accessed January 1, 2026, https://gis-txdot.opendata.arcgis.com/datasets/texas-county-boundaries-detailed/explore.
[4] "Visible Infrared Imaging Radiometer Suite (VIIRS)," NASA Earthdata, accessed January 1, 2026, https://www.earthdata.nasa.gov/data/instruments/viirs.
[5] "VIIRS Nightfire," Earth Observation Group, Payne Institute for Public Policy, Colorado School of Mines, accessed January 1, 2026, https://eogdata.mines.edu/products/vnf/.