Implementing a Quality Program of Analytics Practice: A Framework and Considerations

Michiko I. Wolcott
Msight Analytics, 30725 US Highway 19N #146, Palm Harbor, FL 34684

Abstract
As statisticians, data scientists, and analytics practitioners, we are familiar with the idea of using statistics to measure and improve the quality of other goods and services. Interestingly, relatively little is formally discussed about the quality of our own product: the analytics work. While many analytics practitioners agree the quality of our work should be ensured, articulating what quality even means in our context, let alone ensuring or improving it, can be challenging.

We propose managing analytics quality by applying ideas already recognized in more general contexts, with the operational definition of the quality of analytics practice in terms of defects. We further discuss how this idea translates into a framework for analytics practice across the phases of quality—planning, assurance, control, and maintenance/improvement—and the implications in implementing quality programs for statistics and data science.

Key Words: quality, statistical practice, quality planning, quality assurance, quality control, quality maintenance

1. Introduction

When we think about using analytics (statistics, data science, etc.) to improve the quality, it is to improve the quality of something else. Rarely, if ever, do we think of the quality of the analytics themselves in detail, and when we do, invariably it is either too vague or too narrow. There seems to be a prevailing implicit assumption that analytics practitioners produce high-quality analytics products simply because they are technically capable. It does not matter whether it is predictive modeling, designed experiments, classifications, segmentation, or any other product of statistics, advanced analytics, machine learning, data science, or even AI.

The challenge is that there are so many things that can go wrong that have little to do with technical expertise. Worse, even the most advanced and experienced analytics practitioners are often unaware of the errors they make. I have implemented or have helped implement a quality program in analytics in many organizations over the years, and I cannot understate how common this is.

Analytics practitioners need to be intentional and methodical about how they build quality in their work. For the business or the principal investigator, it is important to know that having analytics done by competent practitioners does not necessarily mean having quality analytics. For the leaders in an organization, it is important to ensure that the analytics function in the organization has a defensible and well-documented quality process.

How is this done? At a recent talk, an attendee poll indicated only about one-third had any quality program or quality methodology for their analytics practice at all. Therefore, the question is rather: how is it that it is not done all the time?

2. Defining "quality" of analytics practice

2.1. Precedents for systematizing analytics quality

There are some precedents for formalizing the idea of the quality of analytics. The traditional approach has been the peer review, which is indeed a method for one piece of it. Since it assesses the analytics product against some good technical practices, it implicitly and qualitatively evaluates the quality of the analytics work.

Others are a little more formalized. The "Quality Assurance Framework of the European Statistical System" (European Statistical System, 2019) is an example in the official statistics realm. In some sectors, there are regulatory mandates with quality implications, like the model risk management requirements in financial services. It is probably reasonable to suspect most of the one-third at that talk had some sort of regulatory requirements, as that has been our observation from experience elsewhere.

There are two major shortcomings of the existing approaches. First, they focus mostly on the analytic output, like precision, rather than on the *practices* that yield quality analytics. They attempt to measure the symptoms of quality rather than to address how to generate quality by design.

Second, the scope is often too narrow for application to analytics practices more generally. Regulatory requirements for the quality of analytics naturally focus on specific aspects. That reduces the scope, not to mention it forces the approach to be more compliance-oriented.

2.2. Defining quality of analytics practice in terms of defects

Most people want what they do to be of good quality—for the pride of workmanship, for reputational reasons, and simply because it is the right thing to do. Quality is also something many people intuitively understand. We all know when it is there. However, quality is also often extremely hard to articulate. Many attempts at analytics quality programs do not go anywhere because they start with what quality notionally looks like, not with what it is.

What is "quality" then? We can start with what the American Society for Quality (American Society for Quality) says:

> "In technical usage, 'quality' can have two meanings:
> 1. the characteristics of a product or service that bear on its ability to satisfy stated or implied needs;
> 2. a product or service free of deficiencies."

The definitions by other established quality experts and organizations are similar. This means it is about meeting some set of criteria or standards, which then leads to the idea of defects. Whenever something does not meet those criteria or standards, it is a defect. This is measurable, as in "the number of defects per million opportunities" often used in a manufacturing setting. Then, the fewer the defects, the higher the quality.

What about the qualitative aspects of quality? We often associate quality with something that we feel we cannot fully measure. However, it still implies there exist some expectations, albeit subjective, for something to be of good quality, since there would not be anything against which to compare our perception otherwise. Whenever that expectation is not met, it can be considered a defect, at least conceptually.

Then, the quality of analytics practice can be defined in terms of defects in analytics practice against some set of standards and expectations.

2.3. The case for the quality of analytics practice

Why should anyone outside of analytics practitioners care about the quality of the analytics practice? Perhaps the most obvious reason is risk management. Poor quality of anything usually has negative consequences, so eliminating defects reduces the associated risks. The most notable risk is that of making a wrong decision from defective analytics. In addition, things like biases and ethical problems can lead to risks that are sometimes not as directly quantifiable.

Second, better quality means more value from analytics. Since each defect is an erosion of value or effectiveness, eliminating defects allows us to get more out of analytics.

Finally, there is the human aspect. Over time, having fewer defects helps develop trust and confidence in the minds of the consumers of analytics. This leads to a greater comfort level with analytics, which is a critical component of adoption and therefore the return on investment (ROI).

2.4. How defects happen in an analytics practice

We all tend to assume we do not make errors. Every time I implement an analytics quality program, the number of defects identified surprises even the most senior analysts. There are two primary reasons for defects in an analytics practice.

The first is competency. Too often, poor-quality analytics result from the analyst not knowing what he or she is doing. It is even more alarming when no one is aware of the poor quality. Unfortunately, it is often discovered only when it causes something to go wrong. While the popularity of everything data has produced an abundance of analysts with the technical knowledge of analytical mechanics, many do not possess a full comprehension of the underlying fundamentals, and this gap can result in devastating defects. That one can apply the techniques does not mean one understands everything involved in using the techniques.

The second is the intent of the analyst. This is the question of whether the analyst meant to do what he or she did. To be clear, an analyst can be fully competent and intentionally do something that somehow does not fit the norm. The worst of these cases is sabotage. However, people do mean well most of the time, and the perceived defect is often a result of specific considerations. In this case, the decision simply needs to be justified and documented; when that does not happen, it then becomes a defect.

The vast majority of the defects happen because competent analysts have not meant what they have done or forgotten why they have done what they have done. It is brought on by carelessness, oversight, inattention to details, lack of diligence, etc., which are innocent but defects, nonetheless. Other things can influence how defects happen. They include whether the analyst means well, whose best interest the analyst has, and whether the analyst even cares, among others.

2.5. The relationship of data quality to the quality of analytics practice

Poor quality of the input data into the analytics process is one of the biggest sources of headaches by analytics practitioners. However, it is important to understand that the input data quality is largely a dependency and a constraint for analytics practitioners. There are some important exceptions, but data quality is not the objective of analytics practitioners.

This is not to say they are not responsible for data quality. Obviously, the quality of the ingredients is a key factor in producing a quality product. Analytics practitioners are indeed responsible for knowing and understanding the quality of the data used. They need to be diligent about studying the quality of the data they use to develop the analytics, to apply appropriate corrections when feasible and to understand and explain the impact data quality has on their product.

However, the role of analytics practitioners is to make something out of data, not to make data error-free. They have to deal with data quality because they depend on it. They need to diagnose data quality issues, apply fixes as appropriate, and explain the limitations of their analysis caused by data quality. However, their primary objective is to draw insights from data, not to improve the quality of the data itself. While analytics practitioners are often at the forefront of data quality, their actions with respect to data quality can only be strictly reactive in the absence of anyone else who is responsible for data management.

Analytics practitioners are not data management professionals. They are not fully versed in all of the data management best (or even standard) practices, and that there is an entire profession dedicated to data management often even comes to them as a surprise. The accountability for the quality of the ingredients belongs—or should belong—elsewhere most of the time, specifically to the chief data officer (CDO) or the equivalent (Msight Analytics, 2018). This is a separate discussion in itself, but a 100% reliance on analytics practitioners to deal with data quality means the bigger data quality problem will never be resolved.

For the business or the researcher, it is important to recognize data quality will remain reactive and unpredictable as long as we expect and depend on analytics practitioners to take care of data quality. Addressing data quality at the organizational level allows analytics practitioners to use their quality mindshare where they should—on activities that produce the analytics.

The point is that analytics practitioners need to focus on what they do to produce a quality product and not get fixated on the circumstances that lead them to produce a quality product. We all tend to be in denial to an extent—it rarely occurs to us to look at ourselves to define quality as a consequence of what we do. As humans, we all tend to look elsewhere for the source of our own improvement. A recent social media post said something to the effect of "the mirror in the bathroom still works." Of course, this is not unique to analytics practitioners, but of all people, analytics practitioners should understand how to build quality in, as W. Edwards Deming long advocated (Deming, 1982).

3. Dimensions of quality for analytics practice

As mentioned earlier, data quality is largely a dependency and a constraint for analytics. However, we can leverage the learnings from industry data quality standards to frame analytics quality. Specifically, it is useful to consider how data management defines dimensions of data quality. Although there are variations, they are fairly well-defined—we can consider, for example, the dimensions as defined in the DAMA Data Management Body of Knowledge (DAMA International, 2017). Data quality issues, in other words, data defects, can be identified against one or more of these dimensions:

- Completeness: How populated or complete is it? Is all required data present?
- Validity: Do the values technically conform to the defined domain?
- Accuracy: Are there errors?
- Consistency: Are the values consistent within and across?
- Integrity: Is the data coherent within and across data objects?
- Timeliness: Does it reflect the timing of interest?
- Uniqueness: Is there any duplication or fragmentation?
- Reasonability: Does it reasonably reflect reality?

We can extend the idea to define a similar set of dimensions for analytics (Figure 1).

Figure 1. Dimensions of analytics quality.

1. Appropriateness: Are the design and the execution of the analysis appropriate for the question and the need?
2. Clarity: Is everything clear and free of ambiguities?
3. Consistency: Is the analysis tool- and system-agnostic? Does it produce the same result every time the same thing is executed?
4. Traceability: Can the lineage of the analysis be traced completely from start to finish?
5. Accuracy: Is the execution free of errors?
6. Transparency: Can the analysis result and the analytic itself be explained? Is everything about the analysis clearly documented so that someone who was not involved can understand?

7. Completeness: Is the logic free of gaps? Has everything that matters been considered and accounted for?
8. Justifiability: Is there a defensible reason for everything? What are the risks, dependencies, and limitations of the choices?

A failure in replicability or reproducibility should have defects in at least one of these dimensions.

Furthermore, it is important to realize the analytics profession does not exist without clients, colleagues, and collaborators. Then, the quality of the analytics work necessarily includes project delivery considerations, since they closely intertwine with the technical aspects of the project (Figure 2):

1. Have clear expectations been set?
2. Have these expectations been met?
3. Has everything, analytical and non-analytical, been justified?
4. Has everything been documented?
5. Is the project free of outstanding issues?
6. Is there validation that all expectations and requirements have been met?

Figure 2: Dimensions of project quality in analytics.

4. Managing the quality of analytics projects

Now that we have the dimensions, we can consider what it means to manage quality in an analytics practice. How do we implement quality practices for analytics?

4.1. The quality management lifecycle

Quality management is a lifecycle, and while there are variations, it generally consists of the following components:

1. Quality planning: defining the requirements, tasks, and activities to design quality into the product.
2. Quality assurance: generating the product by employing methodologies to reduce defects and maximize the likelihood of achieving quality expectations.
3. Quality control: verifying the products meet quality requirements.
4. Quality maintenance and improvement: performing ongoing activities to maintain and improve the quality of the product.

Managing the quality of analytics projects, then, involves applying these principles. It should be noted "quality assurance" and "quality control" are often used interchangeably or even thought to represent the same idea. For our discussion, however, we maintain the distinction above.

4.2. Mapping the quality management lifecycle to analytics

The typical analytics project lifecycle consists of the following stages: discovery, design, development, deployment, and use. The discovery stage is where the opportunity for an analytics project is identified and where the initial discussions take place about the project and the business or the research problem. In the design stage, we plan and design the analytic and the project. The analytic is developed in the development stage then made available in the deployment stage. Finally, the users leverage the analytic made available to them to make business or research decisions in the use stage. Not every practice may have these stages explicitly; however, one can view even a single consultation as a microcosm of this lifecycle with the same concepts.

Figure *3* illustrates how the four components of quality management map against this analytics project lifecycle.

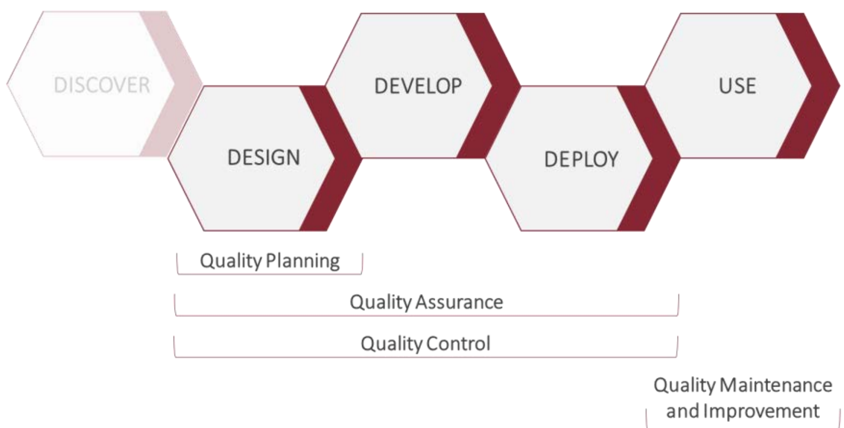

Figure 3: Quality management across the analytics project lifecycle.

Since the discovery stage is typically pre-project, we will set it aside and consider the remaining stages.

- Quality planning takes place primarily in the design stage, in which we define how we plan to design quality into the analysis and the project.
- Quality assurance is carrying out the project and the analysis in a way that minimizes defects and maximizes the likelihood of achieving the quality expectations. Since this applies to practically everything related to an analytics project, it spans from design and development to deployment. There are methodologies, standards, and practices designed to accomplish quality goals. It is also important to standardize approaches, not just routines and macros. Analytics practitioners often shy away from pre-defining, but much more can be standardized than commonly perceived.
- Quality control is verifying that the quality requirements have been met in an analytics project. This often consists of checklists, reviews, and audits. Even design has a doing aspect that can be verified. We do frown upon the concept of

"inspection" in the quality and productivity best practices; however, verification, especially independent verification, that the project meets quality expectations is valuable and sometimes even regulatorily mandated in analytics.

- Quality maintenance and improvement include post-launch or post-publication activities. There are two parts. The first is the ongoing maintenance of the analytic and its use, which primarily concerns (1) ensuring the analytic maintains its external validity over time and (2) identifying operational defects such as system problems and errors, changes in data structures, and changes in contexts and/or behaviors. The second part is improving the system of analytics quality practices for the next project.

4.3. Quality control in analytics practice

Curiously, the traditional notions about quality in analytics have focused on quality control, especially in the form of peer reviews, which may even be regulatorily mandated. This means some quality control practices and standards do exist in analytics, although they may not always be implemented well.

At times, quality control in analytics exists as a system of independent reviews. However, having a separate team dedicated to quality control is not always the best approach. While creating an independent team for this purpose is often encouraged or even required, it does have drawbacks, and it is obviously more challenging for smaller organizations and downright impractical for solo practitioners. There are other ways to implement quality control practices while maintaining some level of independence. The "how" depends on the needs and the circumstances.

It is, however, useful to separate the delivery audit from the expert review within the quality control realm. While the expert review perhaps is more familiar to us, they both question what has been done in the project. **Table *1*** summarizes the key differences between the delivery audit and the expert review.

Table 1: Delivery audit vs. expert review.

	DELIVERY AUDIT	EXPERT REVIEW
OBJECTIVE PROBLEMS	• Expected vs. observed. • Presence of justifications. • Tactical errors (syntax, runtime, semantic). • Proper documentation (analytical and non-analytical).	• Faulty application of analytical expertise. • Recommendations and suggestions based on best technical practices. • Bias and statistical ethics.
STARTING POINT	• "Guilty until proven innocent"	• Observed = Expected. • Everything is justified. • No tactical errors.
SCOPE/FOCUS	• Project delivery	• Statistical/Analytical (i.e., technical)
QUALIFICATIONS	• Lack of expertise can be an asset.	• Specific technical expertise by the reviewer is required.

The purpose of the expert review is to evaluate the analysis against good technical practices. The peer review for journal publications is a classic example. However, many defects in an analytics practice have little to do with technical expertise. To address this, the delivery audit becomes quite literally an audit function. It checks against requirements, ensures analysts have done everything they have said they would, verifies every decision has been justified and documented, and confirms there are no obvious tactical errors.

There are important practical implications of this distinction. While the expert reviewer must be an expert, lacking deep expertise is in fact an asset for the delivery auditor, as innocent questions often lead to error discovery. The expert reviewer should become familiar with the analysis he or she is reviewing, but becoming familiar with the project can compromise the idea of an arm's-length delivery audit.

If we do quality planning and assurance well, then this is straightforward, because we would have already built quality into the product. We "cannot inspect quality into a product" (Deming, 1982), although some verification is always needed, independent or otherwise.

5. Implementing a quality program of analytics practice

It is one thing to understand what quality is and what its theoretical implications are. However, many analytics managers and practitioners struggle with making it a reality, and analytics quality often remains a pipe dream.

So then, how do we make this happen? One thing is for certain: we cannot speak analytics quality into existence. The key practical steps and considerations for implementing an analytics quality program include the following:

- First thing first: create an inventory of the analytics (i.e., the products of analytics work) and the analytics projects. What are we managing the quality of?
- Define roles and responsibilities. Who does what? Should there be an independent team? Who signs off on what? And most importantly, who will own the quality program?
- Define processes and procedures. What does the workflow look like? How should a common flow of activities be defined to make things predictable?
- Define standards, policies, and requirements. What will defects be identified against? Should there be checklists, forms, and auditability standards? Are there any regulatory requirements? The list goes on.
- Set up basic infrastructure. At the minimum, this includes a document store, some logging capabilities, a shared computing environment, and common tools. If we productionalize the resulting analytics, we need appropriate environments, tools, and data sets for testing so that we can test them properly before moving them to production. We have seen from experience that unit testing is inadequate for identifying defects in productionalized analytics.
- Train all analysts on the basic quality principles and practices. They must understand the quality expectations as well as the methodologies and techniques for quality. Since good quality practices often run counter to what analytics practitioners consider elegant and advanced, this requires a shift in mindset.

Obviously, there is a lot more to implementing a successful analytics quality program. It takes balancing the diverse needs of the organization without compromising the quality principles. However, with a solid framework, it is not as specific to techniques or types of analytics as commonly believed. It also requires a lot of discipline, and perhaps not as much art as one might expect.

6. Shifting how analytics practitioners view themselves

This is just the tip of the iceberg.

Analytics practitioners are often resistant to applying statistical principles to themselves. Assumptions are made about the quality of their work, but we seldom check those assumptions as good statistical practices call for. I have come across so many experienced, senior-level analytics practitioners, shocked by all the defects identified upon implementing an analytics quality program. Some become very defensive and start to blame other things like the lack of skills; fortunately, for the vast majority, it is simply a humbling learning experience. A lot of this is human nature—not unlike, as they say, doctors make the worst patients. That said, the analytics practitioners who truly live analytics are invariably the ones most trusted by their clients and colleagues.

It is important to point out the presence of defects does not mean the analyst is incompetent. On the contrary, the vast majority of them are very competent and knowledgeable. It is just that everyone makes mistakes, especially when not being sufficiently diligent. This does require a shift in how analytics practitioners view themselves. A well-respected statistician with decades of experience commented that it had never occurred to him to apply statistical and quality principles to statistics itself. If analytics practitioners are not thinking about how to achieve quality in their analytics systematically, then who is?

References

American Society for Quality. (n.d.). *Quality Glossary - Q.* Retrieved October 5, 2021, from ASQ Quality Resources: https://asq.org/quality-resources/quality-glossary/q

DAMA International. (2017). *DAMA-DMBOK: Data Management Body of Knowledge* (2nd ed.). Basking Ridge, NJ: Technics Publications.

Deming, W. E. (1982). *Out of the Crisis.* Cambridge, MA: Massachusetts Institute of Technology Center for Advanced Educational Services.

European Statistical System. (2019). *Quality Assurance Framework of the European Statistical System.* Retrieved Octoer 5, 2021, from Eurostat: https://ec.europa.eu/eurostat/documents/64157/4392716/ESS-QAF-V1-2final.pdf/bbf5970c-1adf-46c8-afc3-58ce177a0646

Msight Analytics. (2018). *The Oven Engineer's Expertise in a Cookie World: Application to Data and Analytics.* Retrieved October 5, 2021, from Msight Analytics: https://img1.wsimg.com/blobby/go/5118eaea-672f-4570-be36-fc4b60ba4a10/downloads/Msight%20WP%20IM%20Roles%20EN%20201808%20FINAL.pdf?ver=1628545347021