



Trusted Analytics

Mind the gap

It is a common sentiment: as Data and Analytics advances, the techniques and even results are becoming more and more opaque, with analyses operating as 'black boxes'. But for decision makers in data-driven organizations who rely on data scientists and their results, the trustworthiness of these analyses is of the highest importance. In this article, we explore how a concrete approach to Trusted Analytics can help improve trust and throw open the black box of D&A.



drs. M.A.P op het Veld MSc RE
is a partner at KPMG Advisory N.V.
ophetveld.maurice@kpmg.nl



J. Paton MSc
is a Data Science consultant at
KPMG Advisory N.V.
paton.john@kpmg.nl



INTRODUCTION: THE BLACK BOX

Today, complex analytics underpin many important decisions that affect businesses, societies and us as individuals. Biased, gut feel, and subjective decision-making is being replaced by objective, data-driven insights that allow organizations to better serve customers, drive operational efficiencies and manage risks. Yet with so much now riding on the output of Data and Analytics (D&A), significant questions are starting to emerge about the trust that we place in the data, the analytics and the controls that underwrite this new way of making decisions. These questions can be compounded by the air of mystery that surrounds D&A, with increasingly advanced algorithms being viewed by many as an incomprehensible black box.

To address these questions, we will begin this article by examining the current state of trust in analytics among businesses worldwide, to determine if a trust gap exists for D&A. We will then introduce a model for Trusted Analytics, which is a flexible framework for examining trust and identifying areas for improvement in both specific D&A projects and organizations as a whole. We will also examine how our model of Trusted Analytics applies in practice to a selection of D&A activities. Finally, we will conclude by discussing how to close any identified trust gaps in D&A.

DOES A TRUST GAP EXIST FOR DATA & ANALYTICS?

In 2016, KPMG International commissioned Forrester Consulting to examine the status of trust in Data and Analytics by exploring organizations' capabilities across four Anchors of Trust: Quality, Effectiveness, Integrity, and Resilience ([KPMG16]). A total of 2,165 decision makers representing organizations from around the world participated in the survey. Leaders from KPMG, clients and alliance partners also contributed analyses and commentary to this study.

The results of the study were clear. Adoption of D&A is widespread and many companies are clamoring to build their capabilities. Organizations are adopting various types of analytics, from traditional Business Intelligence (BI) to real-time analytics and machine learning. Of the organizations surveyed, at least 70 percent rely on D&A to monitor business performance, drive strategy and change, understand how their products are used, or comply with regulatory requirements. Furthermore, 50 percent say they have adopted some form of predictive analytics and 49 percent say they use advanced visualization, beyond traditional static charts and graphics.

However, a trust gap may hamper relationships between executives and D&A practitioners, and may lead data-derived insights to be treated with suspicion. Only 51 percent of respondents to the survey believe that their C-suite executives fully support their organization's D&A strategy. At the same time, only 43 percent of executives have confidence in the insights they are receiving from D&A for risk and security, 38 percent for customer insights, and only 34 percent for business operations.

Mind the gap

The study also revealed that trust is strongest in the initial data sourcing stage of D&A projects, but falls apart when it comes to implementation and the measurement of the ultimate effectiveness of D&A insights. This means that organizations are unable to attribute the effectiveness of D&A to business outcomes which, in turn, creates a cycle of mistrust that reverberates down into future analytical investments and their perceived returns. Companies which become entrapped in such a cycle run the risk of sacrificing innovative capacity.

Our experience suggests that there are likely several drivers of the trust gap. Decision makers may ...

- ... know that they don't know enough about analytics to feel confident about their use;
- ... be suspicious of the motives or capabilities of internal or external experts;
- ... subconsciously feel that their successful decisions in the past justify a continued use of old sources of data and insight – a form of cognitive bias.

Organizations must think about Trusted Analytics as a strategic way to bridge the gap between decision-makers, data scientists and customers, and deliver sustainable business results

We believe that organizations must think about Trusted Analytics as a strategic way to bridge this gap between decision-makers, data scientists and customers, and deliver sustainable business results.

WHAT IS TRUSTED ANALYTICS?

Most people have a similar instinct for what ‘Trusted Data and Analytics’ means in both their work and their home lives. They want to know that the data and the outputs are correct. They want to make sure their data is being used in a way they understand, by people they trust, for a purpose they approve of and believe is valuable. And they want to know when something is going wrong.

‘Trusted Analytics’ is not a vague concept or theory. At its core are rigorous strategies and processes that aim to maximize trust. Some are well-known but challenging, such as improving data quality and protecting data privacy. Others are relatively new and undefined in the D&A sphere, such as ethics and integrity. We refer to these processes and strategies as the Four Anchors of Trusted Analytics.

1. Quality

Quality is the trust anchor most commonly cited by internal decision-makers. Quality has many dimensions in the D&A space. Key considerations include the appropriateness of data sources; the quality of data sources; the rigor behind the analytics methodologies employed; the methods used to blend multiple data sources together; the consistency of D&A processes and best practices across the organization (and the alignment of these with the wider D&A industry); and the skills and knowledge of data analysts and scientists themselves.

There are many examples of inadvertent quality issues which have had massive knock-on impacts for individuals, organizations, markets and whole economies. And as analytics move into critical areas of society, such as automated recommendation systems for drug prescription, machine learning ‘bots’ as personal assistants and navigation for autonomous vehicles, it seems clear that D&A quality is now a trust anchor for everyone. Most organizations understand and simultaneously struggle with data quality standards for accuracy, completeness and timeliness. As data volumes increase, new uses emerge and regulation grows, the challenge will only increase.

Key concerns in D&A: Quality

- Appropriateness and quality of data sources and data blending
- Rigor and consistency of methodologies and practices
- Skills of data analysts/scientists and alignment with industry best practices and standards

Figure 1. Building trust in analytics ([KPMG16]).



2. Effectiveness

When it comes to D&A, effectiveness is all about real-world performance. It means that the outputs of models work as intended and deliver value to the organization. This is the main concern of those who invest in D&A solutions, both internal and external to the organization. The problem is that D&A effectiveness is becoming increasingly difficult to measure. A reason for this is that D&A is becoming more complex and therefore the ‘distance’ between the upstream investment in people and raw data and the downstream value to the organization is increasing. It is sometimes the case that decision makers do not understand how to evaluate the specific actions being undertaken by analysts and data scientists, or that the greatest impacts of D&A efforts are felt ‘behind the scenes’, e.g. improving access to information or elevating the rigor with which an organization handles its data.

When organizations are not able to assess and measure the effectiveness of their D&A, chances are that decision makers will miss the full value of their investments and assume that a large proportion of their D&A projects 'do not work'. This, in turn, erodes trust and limits long-term investment and innovation. Organizations that are able to assess and validate the effectiveness of their analytics in supporting decision-making can have a huge impact on trust at board level. The result of this, of course, is that organizations that invest without understanding the effectiveness of D&A may not increase the trust or value at all.

Key concerns in D&A: Effectiveness

- Accuracy of models in predicting results
- Appropriate use by employees of D&A insights into their work
- Effectiveness of D&A support in decision-making

3. Integrity

Integrity can be a difficult concept to pin down. In the context of Trusted Analytics, we use the term to refer to the acceptable use of D&A, from compliance with regulations and laws such as data privacy through to less clear issues surrounding the ethical use of D&A such as profiling. This anchor is typically the main concern of consumers and the public in general.

Behind this definition is the principle that with power comes responsibility. Algorithms are becoming more powerful, and can have hidden or unintended consequences. For example, a navigation application could route users past businesses that pay a fee to the developer, or trading algorithms seeking to maximize profit may react unpredictably to unforeseen market circumstances, leading to increased volatility.

How do we decide what is acceptable and what is not? Where exactly does accountability lie, and how far does it reach? This is a new, uncertain and rapidly changing anchor of trust with few globally agreed best practices. Individual views vary widely and there is often no correct answer. Yet integrity has a high media profile and has potentially enormous implications, not only for internal trust in D&A, but also for public trust in the reputation of any organization that gets it wrong.

Key concerns in D&A: Integrity

- Ability to meet regulatory requirements surrounding D&A
- Transparency towards both customers and regulators about data collection and usage
- Alignment with ethical policies and accountabilities

4. Resilience

Resilience in this context is about optimization for the long term in the face of challenges and changes. Failure of this trust anchor undermines all the previous three: it only takes one service outage or one data leak for consumers to quickly move to (what they perceive to be) a more secure competitor. Furthermore, it only takes one big data leak for the regulators to come knocking and for fines to start flying.

Although cyber security is the best-known issue here, resilience is broader than security. For example, many organizations put employees at risk of sharing confidential data with unauthorized people, both inside and outside of their organization. They may further lack controls on which people are permitted to change data. Change management is also very important to this anchor: does the organization follow proven methodologies and practices to enable and take advantage of insights emerging from D&A?

Key concerns in D&A: Resilience

- Tailoring governance policies to specific data use cases
- Authorization and logging for data access, use and analysis
- Cyber assurance for proactively identifying security threats

TRUSTED ANALYTICS IN PRACTICE

The field of D&A is broad and diverse, and the Anchors of Trust can manifest themselves in different forms depending on the context. To examine how Trusted Analytics works in practice, we zoom in to the specialist areas of Business Analytics, Process Mining, and Advanced Analytics.

Business Analytics

Of the three practices that we focus on in this article, Business Analytics has been around the longest and has matured the most. Analytics have been around ever since Enterprise Resource Planning (ERP) systems started to gain ground as the heart of financial administration. The initial rudimentary form of one-off reporting, such as balance sheets and income statements, has over time involved into a whole array of valuable ways to control risks and improve efficiency. With analytics both serving external stakeholders and supporting internal decision-making, consequences of incorrectness are typically financially material. This of course underwrites the necessity of trust.

Trust in Business Analytics is predominantly built on the quality, effectiveness and integrity anchors. Consistency, completeness, correctness and regulatory compliance are key concerns, especially for financial reporting. Organizations look for assurance on the quality of the data they employ by hiring IT auditors to test the data-producing IT infrastructure. Building on the technical integrity of data, the question of reliability emerges. Especially when analytics get more advanced, it is essential that results can be properly interpreted by the intended user. This requires key users to take ownership of the entire process: both upfront, in the functional and technical design, and afterwards, during testing and review. They must also ensure that current and new analytics do not conflict or contradict, unless explicitly designed and properly communicated.

There are several hallmarks of effective Business Analytics. First comes the level to which analytics are embedded in the organization: are analytics available, understood and centrally positioned in the way of working? It is particularly important to stress this for projects that are initiated top-down, where the end-user may not initially recognize the need to use analytics expressed by the board-room. Second comes the alignment between business and IT: the business knows the requirements and will eventually have to use the analytics, IT can expand their understanding of what would be possible.

Integrity and resilience also have their impact on trust in Business Analytics. You could think of compliance with privacy laws: what data about your customers are you preserving, in order to employ fraud analytics? Which analytics do you run on your own employees and how specific do you report? And resilience: in our survey, just 52% of all respondents stated that their data is only changed by authorized people. How do you manage your master data? What governance do you have in place to safeguard resilience?

Example of mishaps in trust

In 2012, US-based retailer Target became the center of a now well-known case of analytics gone wrong. *The New York Times* ([Duh12]) reported that predictive analytics had revealed a teenaged girl's pregnancy, and Target sent her marketing materials geared towards new parents. Unfortunately, the girl's parents were unaware of her pregnancy, and the incident resulted in considerable embarrassment for all parties involved. Target's D&A was clearly of high quality and effectiveness in this case, but their failure to consider integrity still led to a breach in trust.

Process Mining

Like many other disciplines in data science, recent breakthroughs in Process Mining have provided unprecedented opportunities to sense, store, and analyze data in great detail and resolution. Developments in Process Mining have resulted in powerful techniques to discover actual business processes from event logs, to detect deviations from desired process models, and to analyze bottlenecks and waste. However, Process Mining can also be plagued by issues of trust. How can an organization benefit from Process Mining while avoiding trust-related pitfalls?

The quality of a Process Mining analysis can be negatively impacted by both expectations and presentation. Starting from process models may lead to flawed predictions, recommendations and decisions. To provide analysis results with a guaranteed level of accuracy, it is important to use cross-validation techniques to provide adequate confidence about complex analysis results. Resilience can be aided by communicating clearly about the certainty of findings. Where it is necessary, uncertainty about results should be explicitly calculated. At the very least, inconclusive parts of an analysis should be openly presented as requiring further investigation. This can help manage expectations, and is also an opportunity to demonstrate transparency and measure effectiveness. Transparency can be further enhanced by including the ability to drill-down and inspect the data. For example when a bottleneck is detected in a Process Mining project, one needs to be able to drill down to the instances that are delayed due to the bottleneck. It should always be possible to reproduce analysis results from the original data.

Because Process Mining techniques can be used to 'blame' individuals, groups or organizations for deviating from some desired process model, integrity plays

Table 1. The Anchors of Trust: examples of challenges in practice.

	Business Analytics	Process Mining	Advanced Analytics
1. Quality	Accuracy in (financial) reporting	Awareness of initial models and assumptions	Selection of appropriate margins/thresholds
2. Effectiveness	Facilitation of decision making	Aggregation & conformance checking	Asking specific and answerable questions
3. Integrity	Compliance with regulatory requirements	Privacy and discrimination concerns	Ethical issues and privacy concerns
4. Resilience	Data access & modification controls, robustness to change	Scoping and purpose (short-term focus of analyzes)	Agile management and adaptation to ecosystem changes

a key role in trust. Many analysis techniques tend to discriminate among different groups. For example, data analytics can help insurance companies to discriminate between groups that are likely to claim and groups that are less likely to claim insurance. This is often useful and desired, but care must be taken to actively prevent discrimination based on sensitive variables in a given context, for example race, sexuality, or religion. Discriminative-aware Process Mining needs to make a clear separation between likelihood of a violation, its severity, and the blame. Deviations may need to be interpreted differently for different groups of cases and resources.

The Anchors of Trust should be considered in all parts of a Process Mining project including: data extraction, data preparation, data exploration, data transformation, storage and retrieval, computing infrastructures, various types of mining and learning, presentation of explanations and predictions, and exploitation of results taking into account ethical, social, legal, and business aspects.

Example of mishaps in trust

By most accounts, 2016 was a bad year for polling, with pollsters worldwide being consecutively shocked by the outcome of the Brexit referendum in June, and Donald Trump’s victory in the American presidential election in November. Polling agencies use D&A to disseminate information about public opinion, taking great care to ensure integrity by trying to eliminate bias in their methods. However, these recent upsets are very high-profile failures in effectiveness, which in turn lead to questions about quality and harm trust in this application of D&A overall.

Machine Learning / Advanced Analytics

‘Machine Learning’ refers to a set of techniques used to identify patterns in data, without specifically programming which patterns to look for. In this way new insights can be discovered that may not occur to human analysts. Many of these techniques have great potential, but they are almost always statistical by nature: typical results are not precise, but indicative. For example, the results will carry some amount of uncertainty, or they may only apply to a group on average, as opposed to getting specific results for single individuals. This needs to be taken into account when assessing the quality of the output of analyses employing machine learning. Due to the effectiveness of machine learning in identifying complex relationships if they exist in data, the accuracy and therefore quality of predictions combining various sets of data that are available to the organization. The effectiveness of machine learning must be viewed through a similar statistical lens. It can be enhanced by adding rigor to both the design and testing processes, and of course by ensuring that the input data itself is of high quality and consistency.

The integrity of machine learning is a hot topic. With algorithms becoming more complex, it can be difficult to understand what exactly the model has learned, and how it will behave in a certain situation. Ensuring that consumers are treated fairly by algorithms is one of the hallmarks of integrity. Also central to integrity is limiting how much algorithms learn about individual people, to prevent feelings of privacy invasion or ‘creepiness’. Long-term resilience of machine learning can be achieved by resisting the temptation to run many individual exploratory projects without bringing the insights derived into production. A culture of openness can also aid in resilience, helping to shield against (inadvertent) bias creeping into algorithms.

A TRANSITION FROM FINANCIAL AUDIT TO D&A ASSURANCE

Another area where the earlier mentioned concerns are important, is the rise of trust statements or D&A assurance. During internal and external audits, data integrity has always been a key concern and audit focus, this is no different in a data driven economy. However, the need for assurance about data (not necessarily only financial

data), will increase. In a rapidly changing data environment which is impacting how society interacts, the need for an audit (or digital assurance) about the compilation and effect of data driven decision models, will change the audit industry to become an assurance industry. Presently the audit industry is reinventing their business models and service lines in order to find a new fit between the digital revolution in the market and the data and technology transformation that is happening with its clients.

The future of Trusted Analytics: interview with Prof. Dr. Sander Klous

Sander Klous is the Managing Director of Big Data Analytics at KPMG in the Netherlands. He is the founder of this team and is responsible for the global advanced analytics technology stack of KPMG. Sander is professor in Big Data Ecosystems at the University of Amsterdam. He holds a PhD in High Energy Physics and worked on a number of projects at CERN, the world's largest physics institute in Geneva, for 15 years. He contributed to the research of the ATLAS experiment that resulted in the discovery of the Higgs Boson (Nobel Prize 2013). He shares his thoughts on the future of Trusted Analytics.

Q: How have you seen trust in D&A evolving over time?

'Three years ago, when I gave lectures on big data, the biggest question was "What is Big Data?" It was a new concept. Once people started to grasp the fundamentals, they began asking "Now what? How do we get started?", which turned the conversation towards techniques and technology. Then people became concerned with issues like reliability, quantity and quality of data, and how to recruit the right people. These days, when I talk to the board of an organization, they say "D&A is nice, but I'm responsible for the decisions being taken by this organization based on these algorithms, so I need to know that they are reliable (or else I go to jail). But there is another part too: does this algorithm adhere to the norms and values we have as a society? If my algorithm is reliable, but it is always discriminating against a certain race, then it might be reliable and resilient and repeatable, but this is still not acceptable and I will still go to jail." So this is a hot topic these days.'

Q: What is the biggest step that can be taken to improve trust in D&A? Is there a need for a trusted party?

'If you look at the Dutch organization for protecting privacy, APG (Autoriteit Persoonsgegevens), they used to have the obligation to check all the organizations, and if they made a mistake then they had an issue. However, when it came to data leaks, they turned around the responsibility, and made organizations themselves responsible for reporting leaks. This is more scalable, since the work is distributed to all

organizations, and the APG only intervenes in the case of a leak. However, it is clear from the number of reports that organizations are underreporting, so the next step is for authorities to require independent organizations to certify the reporting of data leaks. This is the accountancy model: keep centralized complexity as low as possible, and have the intelligence come from the boundaries, the organizations that are participating in the system. This will first happen for data leaks, but then maybe other regulatory bodies will come with other requirements that need to be checked to ensure that organizations are doing D&A in a responsible manner. So you need independent organizations that know how to deal with trust, and that also have the knowledge to perform these checks.'

Q: Is there a limit to how much control we should give to algorithms and analytics? Should we avoid a situation where all our decisions are made for us?

'Well, we are actually basically in that situation already, and there is even a term for it, from Big Data philosopher Evgeny Morozov: "invisible barbed wire". He says that most of our decisions are already impacted by technology in one way or the other, either voluntarily, for example choosing to use an app for navigation, or involuntarily, for example automatic gates at train stations that deny you access if you have not bought a ticket. The thing is that at the moment, this barbed wire is not really invisible. You still feel it sometimes: when decisions are wrong. If your navigation system directs you down a street that is closed, then you are upset with your navigation system. A decision is made for you and you are not happy with the decision, and that is the barbed wire. Five or ten years from now, the techniques will have improved such that you will not feel the barbed wire anymore, and then it will be too late. We have a window now of five years, maybe ten years, when we can feel the barbed wire, where we as a society can enforce compliance with our norms and values. And just like with the APG and data leaks, this can only be done by reversing the responsibility: making organizations responsible for complying with these values, and subjecting them to regular investigation by a trusted third party. This is the only way to keep the complexity of the task under control.'

CONCLUSION: CLOSING THE GAP

The trust gap cannot be closed by simply investing in better technology. Despite different levels of investment, our survey suggests that more sophisticated D&A tools do little to enhance trust across the analytics life-cycle. We believe that organizations must instead think about Trusted Analytics as a strategic way to bridge the gap between decision-makers, data scientists and customers, and deliver sustainable business results.

In practical terms, this begins with an assessment of the current trust gaps affecting an organization, and a reflection on how to address these. In many cases simple solutions such as implementing checklist-type procedures can already have a big impact. A similar assessment should be conducted for all current D&A activities; this can reveal opportunities for streamlining and alignment.

The quality of D&A can be improved by simplifying interconnected activities as discovered during the initial assessment, encouraging the sharing of algorithm and model design to prevent the (perceived) appearance of 'black boxes', and establishing cross-functional D&A teams or centers of excellence. Cross-functional teams will also help improve effectiveness as they are able to apply their expertise to multiple areas, breaking out of their traditional silos. These teams should be approached with an investor mindset, valuing innovation over the avoidance of failure.

Integrity can be enhanced by fostering a culture of transparency, for example by open-sourcing algorithms and models, and by communicating very openly to consumers about how their data is being used. Openness between business leaders and D&A professionals will also improve resilience, helping to accelerate awareness and align priorities. Resilience is further enhanced by continuously monitoring D&A goals and progress, rigorously test the outcome of analyses and maintaining a whole-ecosystem view of the D&A landscape.

The transition to digital assurance is an opportunity that will have impact on society (people), technology and regulations. A risk emerges if standard setters and regulators would not embrace the digital assurance with the same velocity as the market is doing. This mismatch could develop into a threat for the industry, sparking discussions on relevance of audit procedures in general. A holistic view is vital; each of the parties involved can solve their challenges, but more than before, a multidisciplinary discussion between all parties is required, which may lead to better assurance for corporations and for society in general.

References

- [Duhir12] C. Duhigg, *How Companies Learn Your Secrets*, The New York Times Magazine, February 16, 2012.
- [KPMG16] KPMG, *Building trust in analytics: Breaking the cycle of mistrust in D&A*, KPMG International, 2016.

About the authors

drs. M.A.P op het Veld MSc RE is a partner at KPMG Advisory N.V. He is involved with many audit and advisory engagements for multinationals. At KPMG the Netherlands he is responsible for the data-analytics services. He was editor of *Compact* from 2005 to 2012.

J. Paton MSc is a Data Science consultant at KPMG Advisory N.V. He is a member of the Data and Analytics unit at KPMG, which focuses on advanced analytics training, proof of concepts, and implementation.

Contributions to this article were made by prof. dr. Sander Klous, drs. Mark Kemper, Erik Visser and dr. Elham Ramezani.