Week 7: Data quality, trust, and the harm of bad data at scale

Section 3: Governance

The six dimensions of data quality, documented harm cases, and the accountability gap when the system was wrong.

> **This week's central question.** When a data system produces a wrong answer, someone bears the consequence. In most of the cases on this week's timeline, that someone was not the organization that built or maintained the system. The governance question this week answers is: what accountability structures should exist when bad data causes real harm — and why do the existing ones so consistently fail to produce them?

Data quality is examined this week as both a technical discipline and an ethical obligation. The six dimensions of data quality provide the governance framework; documented cases of high-stakes harm provide the ethics anchor. The cases are not hypotheticals. Robert Williams was arrested in front of his daughters. Davino Watson, a US citizen, spent 1,273 days in immigration detention. Over 40,000 Michigan residents received false fraud determinations with a 93% false positive rate. These are documented outcomes of data quality failures in high-stakes decision systems.

The week's analytical structure follows a chain: which dimension of data quality failed, what decision was driven by the failure, what harm resulted, and who — if anyone — was held accountable. The accountability column in the teaching cases is the week's most important element. In every case it is either empty or inadequate. That pattern is not accidental. It is the product of governance structures that were not designed to produce individual accountability for systemic data quality failures.

The fourth teaching case — Palantir's Investigative Case Management system — is a bridge to Week 9. It introduces the compounding quality problem: when enforcement decisions are made on aggregated data from multiple sources with different quality standards, error rates, and update frequencies, the resulting decision quality is unknown and effectively unauditable. That is the data quality problem at its most dangerous.

LEARNING OBJECTIVES

By the end of this session, students will be able to:

- Define the six dimensions of data quality and identify which dimension is the primary failure in a given case.
- Trace the chain from a specific data quality failure to a downstream harm and identify where in the chain accountability should attach.
- Evaluate the accountability outcomes in the teaching cases and explain why they are structurally inadequate.
- Apply the data justice framework to assess who bears the data quality risk in high-stakes decision systems and who is protected from it.
- Explain why aggregated data from multiple sources compounds the data quality problem and why it makes accountability harder, not easier.
- Connect the data quality discussion to the Week 6 governance frameworks and identify which governance structures are designed to prevent quality failures and which are not.

SESSION STRUCTURE

Hour 1 — The six dimensions: a technical framework with ethical implications

Open with the dimensions table and establish it as the week's diagnostic framework. Walk through each dimension with a brief concrete example, emphasizing the downstream harm: each dimension failure has a characteristic harm profile. Accuracy failures produce wrongful actions against individuals. Completeness failures produce decisions made on partial information. Timeliness failures produce actions based on conditions that no longer exist. The key argument: data quality is not a technical housekeeping concern. It is an ethical obligation when the data drives decisions that affect people's lives, liberty, and livelihoods.

Hour 2 — The harm cases: tracing the chain

Walk through the three anchor cases — Robert Williams, Davino Watson, Michigan UIA — applying the four-part analytical structure to each. Students should be able to name the dimension failure, identify the decision it drove, describe the harm, and evaluate the accountability outcome. The Robert Williams case is the most emotionally immediate — his daughters watching him handcuffed in the driveway is the kind of concrete detail that makes the abstraction real. The Davino Watson case is the most legally consequential — the sovereign immunity outcome means the government owed a citizen nothing for 1,273 days of wrongful detention. The

Michigan UIA case is the most systemic — a 93% false positive rate means the system was wrong more often than it was right.

Hour 3 — The accountability gap: why it is structural, not accidental

Ask students directly: in each of the three cases, who should have been held accountable, and why weren't they? The discussion should surface the structural features that produce the accountability gap: automated systems distribute responsibility across designers, deployers, operators, and maintainers in ways that make individual accountability difficult to assign; sovereign immunity shields government actors from civil liability; class action settlements produce organizational consequences without individual ones; and no professional licensing exists for data governance the way it does for medicine or law. Close with the Palantir ICM case as the bridge to Week 9: when a system aggregates data from dozens of sources with unknown quality, the accountability problem compounds. Who is responsible for the quality of a decision made on data whose quality no one has audited?

Hour 4 — Data contracts and governance responses

Introduce data contracts as a governance mechanism: formal agreements between data producers and data consumers that specify quality standards, SLAs, and accountability for quality failures. Present them not as a technical solution but as a governance instrument that makes quality obligations explicit and assignable. Apply the concept to the teaching cases: what would a data contract between the facial recognition vendor and the Detroit Police Department have required? What would a data contract between Palantir and ICE have required? Students work in pairs to draft the key provisions of a data contract for one of the case organizations, focusing on the accountability clause: what happens when the data quality commitment is breached and a person is harmed?

The six dimensions of data quality

Dimension	Definition	Failure example	Downstream harm
Accuracy	Data correctly represents the real-world entity or event it describes.	A patient record lists the wrong blood type. A criminal database lists the	Wrongful medical treatment; wrongful arrest; wrongful deportation.

Dimension	Definition	Failure example	Downstream harm
		wrong individual for a conviction.	
Completeness	All required data is present; no values needed for the intended use are missing.	A benefits record missing income verification data. An immigration record with no entry for prior legal status.	Benefits denied on incomplete information; enforcement action taken on a partial record.
Consistency	Data is the same across systems and does not contradict itself within or across datasets.	A name spelled differently across three agency systems. A birth date that differs between a passport record and a benefits record.	Identity matching failures; one individual treated as several, or several treated as one.
Timeliness	Data is current enough for its intended use; stale data is not applied to time-sensitive decisions.	An outstanding warrant resolved ten years ago still appears as active. A registered address from a prior decade used for current enforcement.	Individual arrested on a cleared warrant; enforcement action at the wrong address.
Validity	Data conforms to defined formats, ranges, and business rules for its domain.	A date of birth entered as February 30. A zip code with letters. A Social Security Number with nine identical digits.	System errors that produce false matches; data rejected by downstream systems without human review.
Uniqueness	Each entity is represented only once; no duplicates exist that could cause conflation.	Two records for the same individual treated as two different people. Two different individuals whose records are merged into one.	Benefits paid twice or not at all; enforcement action against the wrong individual due to record conflation.

Teaching cases

Robert Williams, Detroit (2020) — wrongful arrest via facial recognition

- **Data quality failure:** accuracy — the facial recognition algorithm misidentified Williams from a grainy surveillance image. The match was flagged as a hit without verification against other evidence.
- **Decision driven by it:** the Detroit Police Department used the facial recognition match as the basis for an arrest warrant without independent corroboration of identity.
- **Harm to the individual:** Williams was arrested in front of his family, held overnight, interrogated for 30 hours, and released when detectives acknowledged the photo did not match. His daughters watched their father handcuffed in the driveway.
- **Who was held accountable:** no officer was disciplined. Detroit PD suspended facial recognition use temporarily, then resumed it. No civil penalty. The ACLU filed suit and settled with policy changes. No individual accountability.

Davino Watson, ICE detention (2008–2011) — wrongful detention via erroneous data match

- **Data quality failure:** accuracy and consistency — Watson, a US citizen, was flagged in an immigration database as undocumented due to an erroneous data match. The error persisted across multiple systems and reviews.
- **Decision driven by it:** ICE detained Watson for over three years based on the database record despite his repeated assertions of citizenship and attempts to obtain his birth certificate.
- **Harm to the individual:** Watson, a US citizen, was detained for 1,273 days. He lost his job, his housing, and years of his life. A federal court awarded him $82,500 — later reduced to zero on sovereign immunity grounds.
- **Who was held accountable:** no individual accountability. The sovereign immunity ruling meant the government owed Watson nothing despite a federal court finding the detention unlawful. The data error that caused it was never publicly explained.

Michigan UIA automated fraud system (2013–2015) — benefits terminated on false fraud determinations

- **Data quality failure:** accuracy and validity — Michigan's MiDAS system flagged unemployment claims as fraudulent based on data mismatches, often minor

discrepancies between employer-reported and claimant-reported wages, without human review.
- **Decision driven by it:** the system automatically issued fraud determinations, assessed penalties of four times the alleged overpayment, and intercepted tax refunds — all without human adjudication of individual cases.
- **Harm to the individual:** over 40,000 people received false fraud determinations. Many lost benefits, had wages garnished, had tax refunds seized, and faced quadruple penalties on amounts they did not owe. Some lost homes. The false positive rate was later found to be 93%.
- **Who was held accountable:** Michigan agreed to a $20 million class action settlement and the system was suspended. No individual accountability for the officials who deployed and maintained it. The vendor faced no consequences.

ICE / Palantir Investigative Case Management (ongoing) — enforcement decisions on aggregated data of uncertain quality

- **Data quality failure:** completeness, accuracy, and consistency failures across aggregated sources — Palantir's ICM ingests data from DMVs, utility records, social media, financial records, and law enforcement databases, each with its own quality standards, update frequency, and error rate.
- **Decision driven by it:** ICE agents use ICM to identify, locate, and prioritize enforcement targets. The system's data quality is not publicly audited; enforcement decisions are made on aggregated profiles whose accuracy is unknown.
- **Harm to the individual:** enforcement actions taken on stale, incorrect, or mismatched data have resulted in documented cases of US citizens detained, individuals arrested at incorrect addresses, and families separated on the basis of erroneous records.
- **Who was held accountable:** no systematic accountability mechanism exists. Individual enforcement errors are litigated case by case. No aggregate audit of ICM data quality has been published. This case connects directly to Week 9's Palantir discussion.

The accountability gap. The accountability outcome in every case above is either empty or describes procedural consequences — policy changes, settlements, suspensions — rather than individual accountability for the people who built, deployed, or maintained the systems that caused the harm. This is not a legal accident. It reflects a structural feature of how automated

decision systems distribute responsibility: the system made the decision, the organization deployed the system, and no individual is clearly accountable for the outcome. That structural feature is the governance problem this week is designed to name.

Seminar discussion questions

1. The Michigan UIA system had a 93% false positive rate — it was wrong more often than it was right. How long should a government agency be allowed to operate a decision system with that error rate before someone is held individually accountable? What governance structure would have caught and acted on that error rate earlier?
2. Davino Watson, a US citizen, was detained for 1,273 days on the basis of an erroneous data match. A federal court found the detention unlawful but sovereign immunity meant the government owed him nothing. Apply the data justice framework: who bore the harm, who was protected from it, and is that distribution ethically defensible?
3. Robert Williams' wrongful arrest resulted in no individual accountability. If the system made the decision, who is responsible for the decision's consequences? Construct the strongest argument for each possible accountability target: the officer who made the arrest, the department that deployed the technology, the vendor who built it, and the officials who approved its use.
4. Data contracts make quality obligations explicit between producers and consumers. What would it mean to extend that concept to the relationship between a government agency and the citizens whose data it processes? What provisions would a "citizen data contract" require — and what enforcement mechanism would make it real?
5. The Palantir ICM system aggregates data from dozens of sources with different quality standards, update frequencies, and error rates, and no aggregate audit has been published. Apply the Week 6 governance failure modes: is this capture, concealment, or bypass — and what governance structure would address it?

Course thread

Coming from — Week 6 examined the governance structures designed to prevent failures. Week 7 examines what happens when those structures fail to prevent harm — and the accountability gap that follows. Students now have both the structural analysis and the human cost.

Going to — Week 8 examines data security and access control as governance infrastructure, connecting the technical controls that prevent unauthorized access to the ethical obligation to protect individuals from the harms documented in Weeks 6 and 7.

Required

Required reading

- ProPublica, “Here’s Where Detroit’s Facial Recognition Went Wrong” (2020) — approximately 15 pages. The primary narrative account of the Robert Williams case.
- Todd Feathers and Alfred Ng, “The Devil in the Details: How Michigan’s Unemployment Agency Became a Predatory Debt Collector” (The Markup, 2021) — approximately 20 pages. The primary investigative account of the MiDAS system and its 93% false positive rate.
- Thomas Redman, Data Quality: The Field Guide (2001) — chapter 1 only, approximately 20 pages. The foundational practitioner treatment of data quality as a discipline; establishes the six dimensions in accessible language.
- Cathy O’Neil, Weapons of Math Destruction (2016) — chapters 1 and 2, approximately 40 pages. The conceptual framework for how automated decision systems produce and obscure harm. Preview of Week 10.

Recommended reading

- Joy Buolamwini and Timnit Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification” (2018) — the technical foundation for understanding facial recognition accuracy disparities. Accessible without a statistics background.
- Angwin, Larson, Mattu, and Kirchner, “Machine Bias” (ProPublica, 2016) — the COMPAS recidivism tool investigation. Preview of Week 10’s algorithmic fairness discussion.
- Michele Gilman, “Poverty Lawgorithms” (Data & Society, 2020) — on automated decision systems in public benefits and the accountability gap for low-income individuals.

Assignment

Enter the course code to unlock assignments.