Print Edition

Saturday 15 August 2026

EDITION NO. 23 · 2 PIECES · 9 PAGES

IN THIS EDITION

Disaster Data's Shaky Foundations Undermine Relief

BY MATHEUS DE SOUZA · IN DEVELOPMENT MAGAZINE · 13 AUGUST 2026

Hurricane Maria hit Puerto Rico in September 2017. In the days following the hurricane, the government reported just 64 deaths. This number reflected only those directly killed by the hurricane — those found the next morning or soon after — and did not include the thousands who died in the weeks and months after the disaster. Maria hit the infrastructure of Puerto Rico hard: there were major power outages that shut down hospitals, medical supplies ran out and other critical infrastructure across the island collapsed. In a special assessment commissioned by the government of Puerto Rico a year later, a team of researchers revised the number of deaths after the event. The total count rose to 2,975 — 46 times higher than the original estimate.

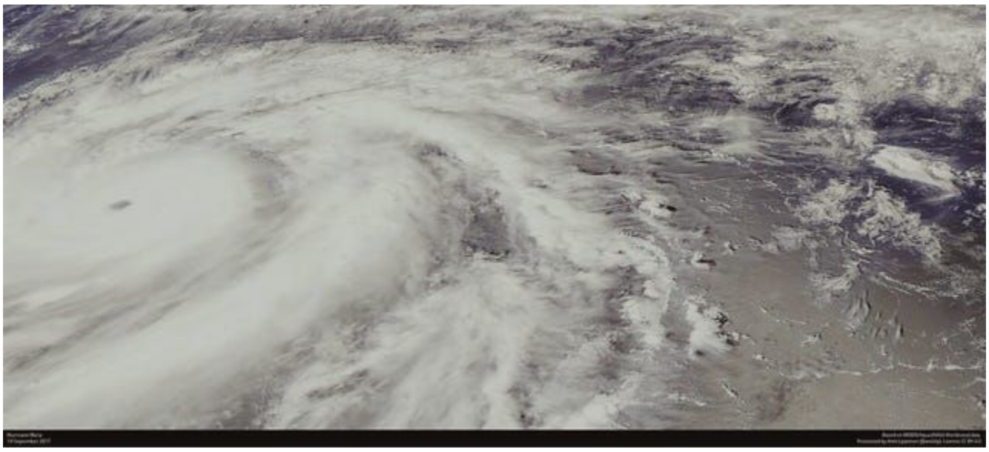

Hurricane Maria from satellite; image by Antti Lipponen, licensed under CC-BY.

The example of Hurricane Maria is not unique. In many countries, the databases used to track the impact of disasters can be badly wrong.

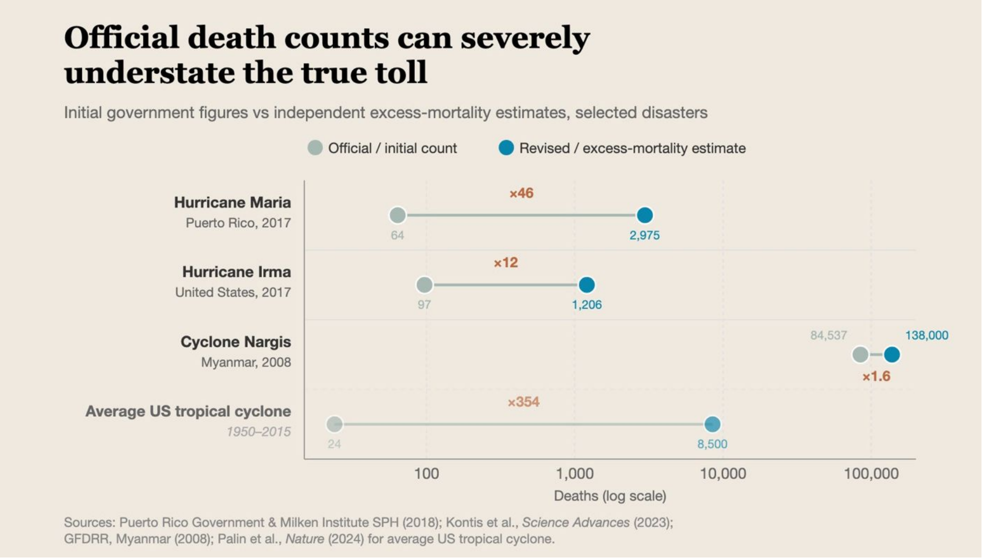

Across contexts, initial death counts underestimate the total death toll.

They rely mainly on government and news reports — but those are shaped by when governments choose to provide information and how these governments define what counts as disaster impact. The data isn't just used for academic purposes; disaster figures weigh heavily in today's world. They can influence which countries receive funding and where disaster funds are spent.1

While these decisions are often framed as technical, there is a political dimension to disaster reporting. The data can often determine whose suffering is recorded and prioritized. We treat these figures as facts, but they are not. They are generated by systems, with the same fragilities as any other policy realm. Disaster data is far less reliable than its users often assume.

The criteria for good data are not controversial. Data should reflect what actually happened, capture the full picture, remain comparable across time and geography, and not count the same event twice. Unfortunately, we are far from achieving that.

How Disaster Databases Actually Work

We need data from disasters. Recovery efforts must be planned, humanitarian assistance delivered and resilience plans made. The Emergency Events Database (EM-DAT) is the most widely used free source of global disaster data, covering both technological and natural hazards. It has records of more than 27,000 events since 1900. For an event to be included, it must meet at least one of three main criteria: 10 or more people killed; 100 or more people affected; or a declaration of a state of emergency or call for international assistance.

While reasonable, these criteria mean the dataset has blind spots. A drought that killed nine people will typically not be included in the database, unless it triggers an emergency declaration or an international call for aid. The same occurs for a flood that affects 95 people. And there's an important caveat there: it is focused on the number of people reported affected or killed. If the reporting isn't accurate, neither is the database. EM-DAT requires cross-verification of each event from at least two independent sources, which are often international agencies, wire services and English-language media.

Given this, geographical bias in media reporting is likely to propagate into the EM-DAT database itself, leading it to capture more disaster events in developed economies than in less developed ones. It also lets governments manipulate the data. If the government only admits to nine deaths, the event won't be included — even if there are far more than nine deaths in reality. This lets governments evade accountability for their actions.

Consider the 2008 Sichuan earthquake in China. Several of the hardest-hit areas were poorer counties in the region, and many buildings collapsed. The official death count was fixed at around 70,000, with a further 18,000 people still listed as missing. Parents, activists and journalists who tried to seek accountability and investigate construction codes and the collapse of buildings were targets of censorship, detention and surveillance; the full toll and responsibility, especially for the school collapses, remain difficult to verify independently.

Even when an event meets the thresholds, other information may be incomplete. Information on economic losses is the main source of missingness; it is unavailable for 80% of the events recorded from 2000 to 2020. Economic damages are hard to calculate and are rarely provided in low- and middle-income countries. Just 4% of African disasters recorded in EM-DAT have economic damage estimates.

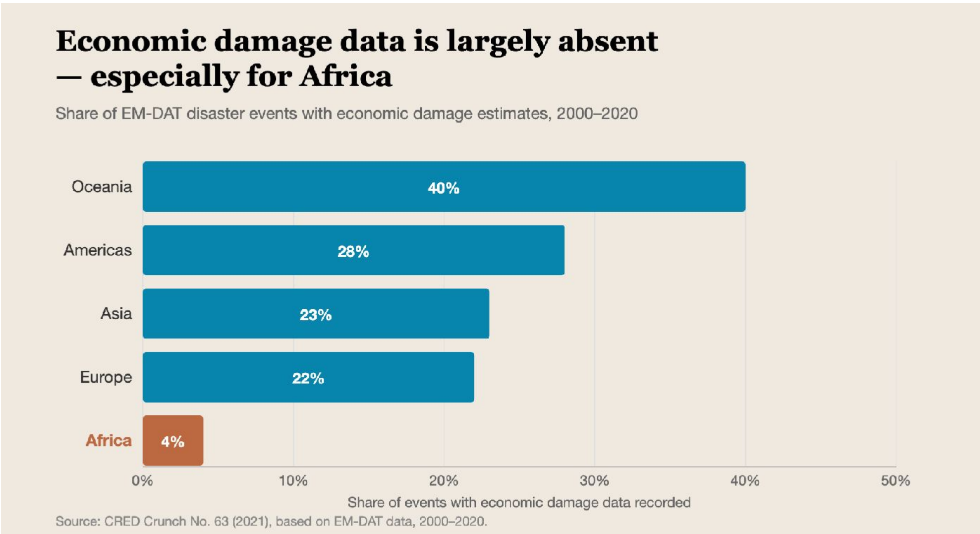

While no region has complete coverage of economic damages, Africa has particularly bad coverage.

But this doesn't stop people from using the data off-the-shelf. A growing number of highly influential and widely cited publications use this data for empirical studies with very limited acknowledgement of its problems.

The scale of the limitations becomes clearer when you compare disaster databases. Between 1971 and 2002, EM-DAT recorded 97 disasters in Colombia. Another database, DesInventar, recorded more than 19,000 in the same period. Only a small percentage of the local events in DesInventar would meet EM-DAT's definition of a disaster, but taken together, the "small" events in Colombia caused more than $1.65 billion in damages. This is around seven times the economic losses caused by one of Colombia's deadliest disasters, the Nevado del Ruiz volcanic eruption. Together, these "small" events pack a big punch.

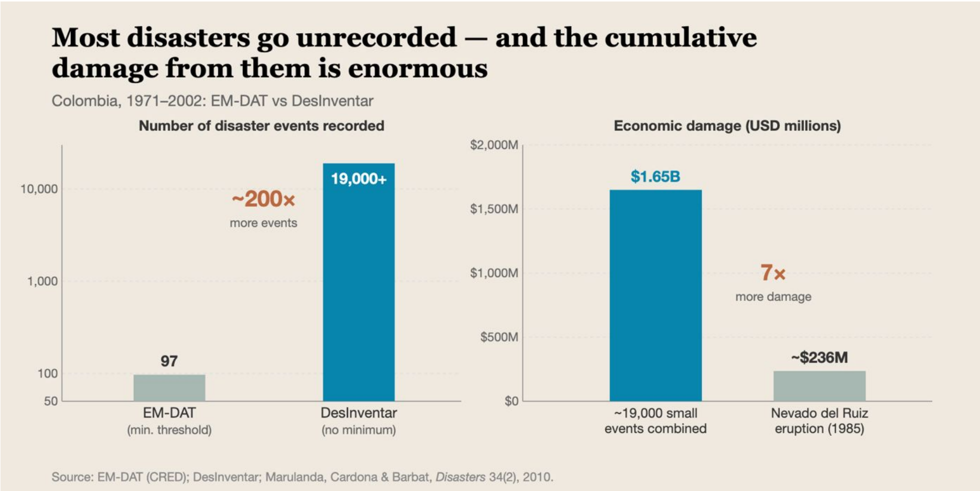

There are many small disasters; together, they cause as much damage as large events.

DesInventar represents a different type of disaster database. Instead of setting specific criteria for inclusion, DesInventar allows countries to freely build their own disaster inventories. In general, this means that it includes smaller disaster events, but it also means the quality of the data included within the database depends on each government's capacity to collect and maintain records. This can vary considerably. Peru is a prime example. Its data uses a single source — and events near Lima are far better covered than events in more remote parts of the country.

Nor is that the only discrepancy. The count of those who are "affected" can vary substantially. EM-DAT defines those affected as requiring immediate assistance, but countries have different standards for what is considered "assistance." Some include all those in the disaster zone; others include only those who lost housing, while still others count assistance requests as the parameter. Countries might even strategically choose who counts as affected, because such data can help determine if a country qualifies for climate adaptation financing or not.

Other categories of data are even murkier. As noted, economic damage estimates are often missing, and where they do exist, the data is sketchy at best. Some studies comparing disaster events between EM-DAT and DesInventar found that damage figures diverged by more than 20% in *most* cases.

Turning Disaster Data Into Policy

This may seem an academic concern. It is not. These databases are used in real policy decisions.

The United Nations Office for Disaster Risk Reduction uses both EM-DAT and DesInventar for its Global Assessment Reports, one of the main publications guiding international disaster policy. United Nations Member States also report their progress through the Sendai Framework Monitor, which uses disaster databases as inputs. This is then used to track progress, identify disaster-risk priorities, guide resource allocation toward risk reduction and help plan climate adaptation strategies.

Climate finance, a sector that moved $1.9 trillion in 2023 alone, also uses disaster databases to assess climate-related disaster risk. Countries often use Germanwatch's Global Climate Risk Index to discuss the need for funding support for adaptation or losses. Germanwatch's current index uses EM-DAT with few caveats on its limitations. This loss and damage analysis, for instance, uses the index without discussing its data limitations.

The reliance on EM-DAT goes beyond the Global Climate Risk Index and into the operational structures that determine humanitarian funding flows. The INFORM Risk Index uses EM-DAT for a few of its hazard indicators and has been used by the United Nations Office for the Coordination of Humanitarian Affairs, the European Commission's Civil Protection and Humanitarian Aid Operations, the World Food Programme, the United States Agency for International Development, the UN Central Emergency Response Fund and more to help prioritize and allocate humanitarian resources. Mistakes in EM-DAT could leave disasters underfunded. This might occur if many more people died than are reported in EM-DAT because their deaths were not reported in English.

In 2015, the Malawi government paid $5 million for drought insurance through the African Risk Capacity (ARC), which pays out when its model estimates that a drought has crossed a given threshold. After the 2016 El Niño drought, which left 6.5 million people in need of aid, the model initially found that no payout was warranted, based on a false assumption regarding the maize variety farmers had planted. A payout was made in 2017, by which point the response had cost around $395 million.

Building Firmer Ground: Five Fixes

We clearly need disaster data. Without databases like EM-DAT and DesInventar, we would not be able to understand and respond to disasters. It would be far more difficult to respond to climate change and plan adaptation strategies to prevent future disasters. But they are far from perfect.

Improvement is possible, though. Indeed, as the world warms, and climate-related disasters become more frequent, improvement is necessary.

First: we must begin to treat disaster databases not as a single source of truth but as part of a system. Each disaster database is a measurement tool with reported limitations, but considering them together can limit the errors. Resource allocation should be triangulated between different disaster databases and local disaster inventories. This will help limit missing data — especially vital given that smaller, recurring events represent an important share of the disaster impact.

Secondly, disaster impact should be reported in ranges, not in a headline figure alone. It is too easy for policymakers to focus on a single number without taking into account the uncertainty behind it. In 2015, the Integrated Research on Disaster Risk program recommended that disaster data should include reliability information such as a quality score or uncertainty level. Initially reported mortality figures from Hurricane Maria were 46 times lower than those produced by the excess mortality analysis.

A reporting range would not have eliminated the discrepancy, but it would signal to policymakers that the number was provisional rather than definitive.

Third, we must verify official counts independently. Governments are not always incentivized to tell the truth, and even if they are attempting to do so, different standards can lead to vastly different estimates of damage. Instead of relying on a count of people in a damaged area, we should focus on excess mortality as a standard post-disaster metric. This would provide an independent data point against which official figures can be assessed; if a government gives a figure that is many times less than the number of excess deaths, they can be called to account.

We can also use satellite imagery to validate data. This was used to great effect after the Haiti earthquake in 2010, as well as in Turkey after the 2023 earthquakes. These technologies cannot replace ground-level data collection, as they cannot count individual people. However, they provide a form of independent analysis that can be used to compare against other figures that enter disaster databases. This is particularly useful in contexts with weak data collection or reporting infrastructures; in countries like Haiti, it can be nearly impossible to collect data door-to-door.

The fourth fix goes hand in hand with this. We must also better fund local data systems. Several countries have adopted DesInventar-based systems for their own disaster loss accounting, but there are few incentives for maintaining these systems. Every year, governments need to train personnel, validate data and verify that interconnected systems actually function together. Each of these costs money, and none happens automatically; international organizations should help support this where possible.

Fifth, make the limitations of the data visible. No policy document should use disaster data without caveats. Each disaster database comes with its own biases; none is perfect. Without caveats, policymakers may take the data off-the-shelf without considering those biases. Every report that cites disaster data should include material like SDG monitoring metadata, INFORM's reliability score and Climate Analytics' loss and damage briefing to provide the reader with context about the data reliability.

No system is perfect, and the systems built to account for the impacts of disasters, including who was affected, the size of the human losses and the economic cost, will always carry intrinsic assumptions, political pressures and institutional limitations. But we can do better than the current system, particularly as natural disaster frequency increases due to climate change. Preventable mistakes will come with a death toll, and right now, disaster databases risk mistaking what is currently reported for what matters.

Matheus de Souza is a PhD student in Disaster Science and Management at the University of Delaware and has worked with different humanitarian organizations such as the United Nations High Commissioner for Refugees and the Norwegian Refugee Council. He studies several aspects of disasters, with a focus on inter-organization coordination.

If you have comments on this article, or wish to contribute to the discussion, please email them to letters@indevelopmentmag.com. Responses will be featured in a letters section.

1

For instance, the UN Central Emergency Response Fund uses the INFORM Risk Index as part of its analysis to identify underfunded emergency allocations and recommend the size of their funding.

AI, Productivity, and Rates: Part I

BY MATTHEW C. KLEIN · THE OVERSHOOT · 5 AUGUST 2026

Two videos you may be interested in from my recent trip to Dalian: a panel discussion I did on "Resisting Autarky", and a free-wheeling conversation I had with Nikhil Kamath and Zhu Ning.

According to the Federal Reserve's latest policy statement, "productivity growth and capital investment are strong". At the press conference following the latest meeting, Fed boss Kevin Warsh elaborated that the U.S. economy was in “a race” between efficiency gains, which allow workers to produce more goods and services from the same volume of inputs, and business spending, which has been driving up the prices of everything from land to laptops. Both, of course, are being driven by AI. The outcome of this race, according to Warsh, will determine what happens to inflation, and, by extension, the appropriate level of interest rates.

There is no doubt that American companies have been spending hand-over-fist on equipment for data centers, lately to the chagrin of shareholders and creditors. Productivity, however, is much harder to measure. Output per hour is estimated by dividing three separate data series—nominal spending and incomes, prices, and hours worked—all of which are prone to measurement error and revision.

That said, the data we have so far from the Bureau of Economic Analysis (BEA) and the Bureau of Labor Statistics (BLS) imply that American workers produced about 2% more goods and services per hour worked in 2026Q2 than in 20252. That would be only slightly above the longer-term trend of productivity growth from 2004Q1-2019Q4, which also happened to be the same as the trend growth rate of productivity from 1984Q1-1996Q4.[1] (Things looked slightly better a few quarters ago.) So far, the excitement about productivity is more about expectations than what is currently visible in the hard data.

Breaking the Trend?

Current estimates of output per hour worked are only slightly above what one would have expected based on prior trends. If productivity accelerates, it could affect interest rates.

Nonfarm business output per hour worked* (1970Q1=100)

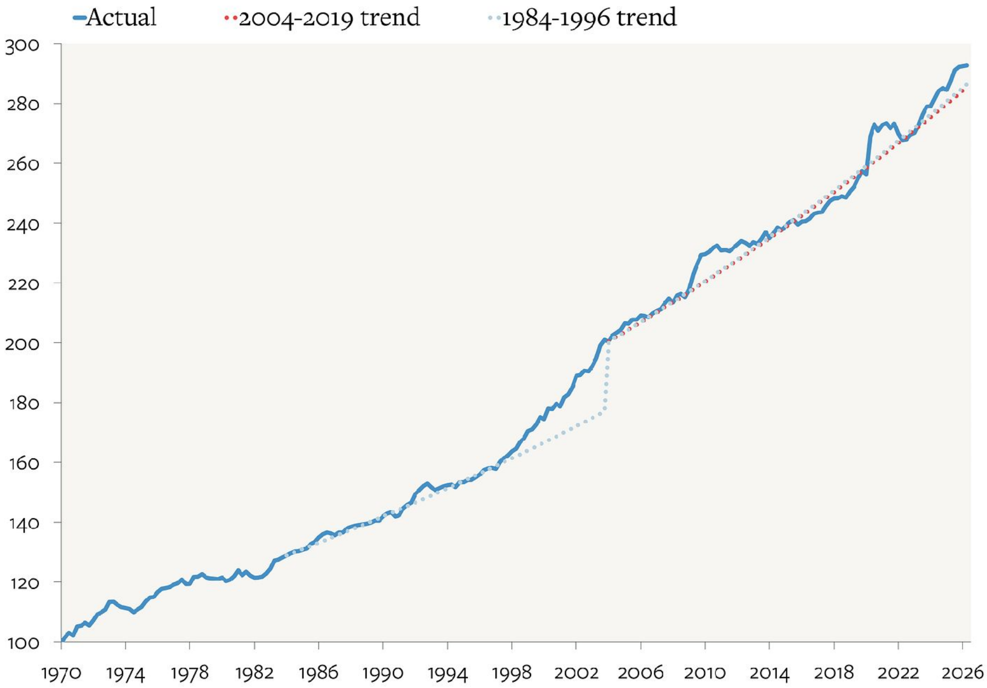

*2026Q2 value estimated based on weekly hours worked by employees on nonfarm private payrolls

Source: Bureau of Labor Statistics, Bureau of Economic Analysis, Matthew Klein’s calculations

But even if productivity growth *were* poised to accelerate, and even if that coincided with an "immaculate" disinflation, which so far seems unlikely, the implications for interest rates are far from obvious. The lesson of the 1990s, to the extent that there is one, is that real rates might need to increase by *more* than any drop in inflation compensation, depending on what new technologies do to the prospective returns on investments and the resulting demand for capital.

The danger is that policymakers inadvertantly lean into a boom, thereby turning it into a bubble that then bursts violently. The same animal spirits and financial leverage that inflate activity on the way up are incredibly destructive on the way down. It would be far better for policymakers to focus on smoothing out the cycle—which is, after all, their job—by stretching out the investment boom over time. That could require higher interest rates even if inflation somehow manages to decelerate to the Fed's alleged 2% yearly target.

While writing this note, I realized that I needed to split it into two parts. This one focuses on the U.S. productivity data and why new software has not yet had much of an impact on standard measures of technological progress. The second part will focus on the implications for interest rates and monetary policy, with a deeper look at the experience of late 1990s and early 2000s.

Is Productivity Accelerating or Slowing?

While the headline measure of output/hour had been rising relatively briskly, there are reasons to be skeptical about how sustainable this is, much less whether it is attributable to genuine innovation.

Output/hour can rise for many reasons. Workers can get more experience and education, they can get equipped with more and better machines, they can benefit from cheaper input prices—someone on the other side of the world starts pumping more oil, for example—and they can push themselves to work harder in ways that might not be immediately evident in the data. While all of these changes can be good for growth, none of them make workers more efficient, because they all involve using more inputs to generate more outputs.

This is why economists at the Federal Reserve Bank of San Francisco estimate the contribution of underlying improvements in technology and efficiency, which they call "utilization-adjusted total factor productivity" (TFP), by stripping out the growth impact of changes in hours worked, the capital stock, "labor quality", and the state of the business cycle. (Those who want the details should read the technical paper by John Fernald explaining the methodology.)

The chart below shows how all of these factors have contributed to quarterly changes in "real business output", which is close to, but conceptually distinct from, both Gross Domestic Product (GDP) and the measure (found in NIPA table 1.3.3.) used by the BLS in their official series on productivity.2

CONTINUES ONLINE · https://theovershoot.co/p/ai-productivity-and-rates-part-i