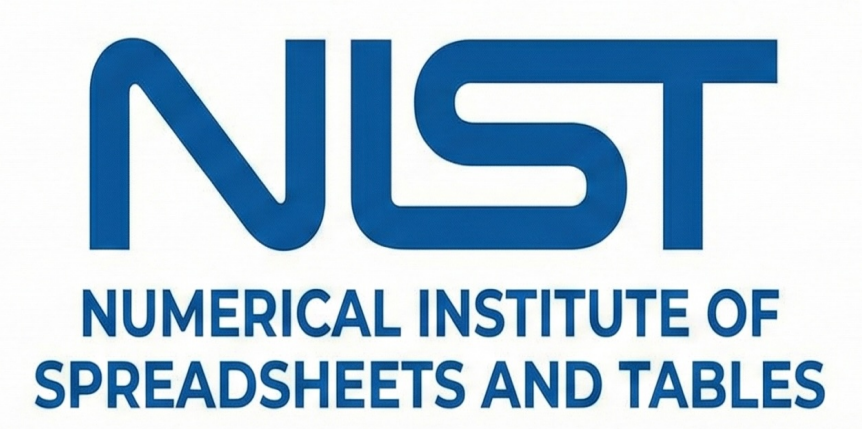

NIST SS 100-1

Spreadsheet Risk Management Framework (SRMF 1.0)

"Because every VLOOKUP is a policy decision."

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SS.100-1
(This is not a real URL)

December 2025

DISCLAIMER: This document is a satirical work created for educational and entertainment purposes. It is not affiliated with the National Institute of Standards and Technology (NIST). However, the spreadsheet incidents described herein are real, and the framework structure parodies the actual NIST AI RMF 1.0.

NIST SS 100-1

Spreadsheet Risk Management Framework (SRMF 1.0)

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SS.100-1
(This is not a real URL)

December 2025

U.S. Department of Tabular Commerce
Someone Important, Secretary

Numerical Institute of Spreadsheets and Tables
A Person Who Uses Excel, Director

Table of Contents

Executive Summary

Spreadsheet technologies have significant potential to transform society and people's lives – from commerce and budgeting to healthcare administration to economic policy modeling to tracking the movements of persons of interest by intelligence agencies. Spreadsheet technologies can drive inclusive economic growth and support administrative advancements that improve organizational efficiency. Spreadsheet technologies, however, also pose risks that can negatively impact individuals, groups, organizations, communities, society, the environment, and the planet. Like risks for other types of technology, spreadsheet risks can emerge in a variety of ways and can be characterized as long- or short-term, high- or low-probability, systemic or localized, and high- or low-impact.

> The SRMF refers to a **spreadsheet system** as a tabular data structure, typically consisting of rows and columns containing formulas, values, and references, that can generate outputs such as calculations, projections, or reports influencing real or virtual environments. Spreadsheet systems are designed to operate with varying levels of complexity (from a single SUM formula to models determining national economic policy). (Adapted from: Someone's best guess; Cell A1 of our founding document)

While there are myriad standards and best practices to help organizations mitigate the risks of traditional software or information-based systems, the risks posed by spreadsheet systems are in many ways unique (See Appendix B). Spreadsheet systems, for example, may contain formulas that reference data that can change over time, sometimes significantly and unexpectedly, affecting system functionality and trustworthiness in ways that are hard to understand. Spreadsheet systems and the contexts in which they are deployed are frequently complex, making it difficult to detect and respond to failures when they occur. Spreadsheet systems are inherently socio-technical in nature, meaning they are influenced by societal dynamics and human behavior. Spreadsheet risks – and benefits – can emerge from the interplay of technical aspects combined with societal factors related to how a system is used, its interactions with other spreadsheet systems, who operates it, and the social context in which it is deployed.

These risks make spreadsheets a uniquely challenging technology to deploy and utilize both for organizations and within society. Without proper controls, spreadsheet systems can amplify, perpetuate, or exacerbate inequitable or undesirable outcomes for individuals and communities. With proper controls, spreadsheet systems can mitigate and manage inequitable outcomes. However, we acknowledge that establishing such controls is difficult, expensive, and frequently ignored.

Spreadsheet risk management is a key component of responsible development and use of spreadsheet systems. Responsible spreadsheet practices can help align the decisions about spreadsheet system design, development, and uses with intended aims and values. Core concepts in responsible spreadsheet use emphasize human centricity, social responsibility, and sustainability. Spreadsheet risk management can drive responsible uses and practices by prompting organizations and their internal teams who design, develop, and deploy spreadsheets to think more critically about context and potential or unexpected negative and positive impacts. Understanding and managing the risks of spreadsheet systems will help to enhance trustworthiness, and in turn, cultivate public trust. Or at minimum, reduce the frequency of multi-billion dollar losses.

> **Responsible spreadsheet use** is meant to result in technology that is also equitable and accountable. The expectation is that organizational practices are carried out in accord with 'professional responsibility,' an approach that aims to ensure that professionals who design, develop, or deploy spreadsheet systems recognize their unique position to exert influence on people, society, and the future of tabular data. This includes the recognition that 'it works on my machine' is not a validation strategy.

The Framework is divided into two parts. Part 1 discusses how organizations can frame the risks related to spreadsheets and describes the intended audience. Next, spreadsheet risks and trustworthiness are analyzed, outlining the characteristics of trustworthy spreadsheet systems, which include valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy enhanced, and fair with their harmful biases managed.

Part 2 comprises the 'Core' of the Framework. It describes four specific functions to help organizations address the risks of spreadsheet systems in practice. These functions – GOVERN, MAP, MEASURE, and MANAGE – are broken down further into categories and subcategories. While GOVERN applies to all stages of organizations' spreadsheet risk management processes and procedures, the MAP, MEASURE, and MANAGE functions can be applied in spreadsheet system-specific contexts and at specific stages of the spreadsheet lifecycle.

Part 1: Foundational Information

1. Framing Risk

Spreadsheet risk management offers a path to minimize potential negative impacts of spreadsheet systems, such as threats to financial stability, public health tracking, and organizational decision-making, while also providing opportunities to maximize positive impacts. Addressing, documenting, and managing spreadsheet risks and potential negative impacts effectively can lead to more trustworthy spreadsheet systems.

1.1 Understanding and Addressing Risks, Impacts, and Harms

In the context of the SRMF, risk refers to the composite measure of an event's probability of occurring and the magnitude or degree of the consequences of the corresponding event. The impacts, or consequences, of spreadsheet systems can be positive, negative, or both and can result in opportunities or threats. When considering the negative impact of a potential event, risk is a function of 1) the negative impact, or magnitude of harm, that would arise if the circumstance or event occurs and 2) the likelihood of occurrence. Given that studies suggest 88-94% of spreadsheets contain at least one error, the likelihood of occurrence approaches certainty for any spreadsheet of sufficient complexity.

> "Risk management refers to coordinated activities to direct and control an organization with regard to risk" (Source: ISO 31000:2018). In the spreadsheet context, this primarily involves hoping that nobody notices the errors.

The SRMF is designed to address new risks as they emerge. This flexibility is particularly important where impacts are not easily foreseeable and applications are evolving. While some spreadsheet risks and benefits are well-known, it can be challenging to assess negative impacts and the degree of harms. The following examples illustrate the range of potential harms related to spreadsheet systems:

Harm to People	Harm to an Organization	Harm to an Ecosystem
• Individual: Harm to a person's benefits determination, credit assessment, or employment decision based on erroneous spreadsheet output	• Harm to an organization's financial position (see: JPMorgan $6.2B)	• Harm to interconnected systems relying on spreadsheet outputs
• Group/Community: Harm to populations through flawed policy analysis (see: Reinhart-Rogoff)	• Harm from public health data loss (see: PHE COVID-19)	• Harm to scientific research integrity (see: gene name crisis)
• Societal: Harm to democratic processes through flawed election tabulation	• Harm to reputation when errors become public	• Harm to regulatory systems and oversight

Fig. 1. Examples of potential harms related to spreadsheet systems. Trustworthy spreadsheet systems and their responsible use can mitigate negative risks and contribute to benefits for people, organizations, and ecosystems.

1.2 Challenges for Spreadsheet Risk Management

Several challenges are described below. They should be taken into account when managing risks in pursuit of spreadsheet trustworthiness.

1.2.1 Risk Measurement

Spreadsheet risks or failures that are not well-defined or adequately understood are difficult to measure quantitatively or qualitatively. The inability to appropriately measure spreadsheet risks does not imply that a spreadsheet system necessarily poses either a high or low risk. In fact, the most dangerous spreadsheet is often the one that appears to work correctly. Some risk measurement challenges include:

Hidden complexity: The most dangerous formula is the one in a cell you didn't know contained a formula. Unlike traditional software, where code is visible in dedicated files, spreadsheet logic is distributed across thousands of cells, any of which may contain formulas, hardcoded values, or references to external sources. The visual presentation of a value gives no indication of whether that value is computed, static, or linked.

Silent failure modes: Unlike traditional software, spreadsheets fail silently. A broken VLOOKUP returns #N/A, not a stack trace. A formula that references the wrong range continues to produce plausible-looking numbers. An excluded row in a SUM still yields a total. These silent failures can propagate through interconnected calculations without triggering any visible error state.

Institutional knowledge decay: The person who built the spreadsheet has always left the organization. This is not an exaggeration; it is a statistical inevitability given employee turnover rates and the typical lifespan of organizational spreadsheets. Documentation, if it exists, is rarely updated to reflect subsequent modifications.

Version proliferation: 'Budget_FINAL_v3_REAL_USE_THIS_ONE(2).xlsx' represents a failure mode unique to spreadsheet systems. The ease of copying and modifying spreadsheets, combined with inadequate version control practices, leads to multiple divergent versions with unclear provenance and authority.

1.2.2 Risk Tolerance

While the SRMF can be used to prioritize risk, it does not prescribe risk tolerance. Risk tolerance refers to the organization's readiness to bear the risk in order to achieve its objectives. In practice, organizations demonstrate remarkably high risk tolerance for spreadsheet errors, often because they are unaware of the errors' existence. Risk tolerances are likely to change dramatically after a public incident.

> The Framework is intended to be flexible and to augment existing risk practices which should align with applicable laws, regulations, and norms. Organizations should follow existing regulations and guidelines for risk criteria, tolerance, and response established by organizational, domain, discipline, sector, or professional requirements. Where established guidelines do not exist (which is most places for spreadsheets), organizations should define reasonable risk tolerance. Once tolerance is defined, this SRMF can be used to manage risks and to document risk management processes. Also, consider whether the thing should be a spreadsheet at all.

1.2.3 Risk Prioritization

Attempting to eliminate negative risk entirely can be counterproductive in practice because not all incidents and failures can be eliminated. Most spreadsheets contain errors; most errors are never detected; and the resources required to achieve error-free spreadsheets exceed the resources most organizations are willing to allocate. A risk management culture can help organizations recognize that not all spreadsheet risks are the same, and resources can be allocated purposefully.

When applying the SRMF, risks which the organization determines to be highest for the spreadsheet systems within a given context of use call for the most urgent prioritization and most thorough risk management process. In cases where a spreadsheet system presents unacceptable negative risk levels – such as where significant negative impacts are imminent, severe harms are actually occurring, or catastrophic risks are present – deployment should cease until risks can be sufficiently managed. If a spreadsheet system's development and use cases are found to be low-risk in a specific context, that may suggest potentially lower prioritization. Note that most spreadsheets are categorized as 'Low' until an error is discovered, at which point they are retroactively understood to have been 'High' all along.

2. Audience

Identifying and managing spreadsheet risks and potential impacts – both positive and negative – requires a broad set of perspectives and actors across the spreadsheet lifecycle. Ideally, spreadsheet actors will represent a diversity of experience, expertise, and backgrounds. The SRMF is intended to be used by spreadsheet actors across the spreadsheet lifecycle and dimensions.

Spreadsheet actors displayed in this framework include:

- **Spreadsheet Developers:** Those who create or modify spreadsheets, whether that is their job title or not (it is rarely their job title). These actors are responsible for formula logic, data structures, and the initial design decisions that shape system behavior.
- **Spreadsheet Deployers:** Those who distribute spreadsheets for use by others, often via email attachment with no documentation. These actors are responsible for decisions about how and where the spreadsheet will be used.
- **Spreadsheet Consumers:** Those who use spreadsheet outputs to make decisions, typically without examining the underlying logic. These actors rely on the accuracy of outputs they cannot verify.
- **Spreadsheet Governors:** Those responsible for organizational policies regarding tabular data. This role exists in approximately 0.3% of organizations. Where it does exist, it is typically assigned to whoever made the mistake of demonstrating competence with Excel.
- **Spreadsheet Inheritors:** Those who must maintain spreadsheets they did not create, for purposes they do not fully understand, using formulas they cannot interpret. This is the most common role.
- **Affected Individuals:** Those whose benefits, employment, credit, or other life outcomes are determined by spreadsheet calculations they may never see or know about.

Successful risk management depends upon a sense of collective responsibility among spreadsheet actors. The SRMF functions require diverse perspectives, disciplines, professions, and experiences. Diverse teams contribute to more open sharing of ideas and assumptions about the purposes and functions of technology – making these implicit aspects more explicit. This broader collective perspective creates opportunities for surfacing problems and identifying existing and emergent risks. Unfortunately, spreadsheets are often created by a single person working in isolation, which undermines these benefits.

3. Spreadsheet Risks and Trustworthiness

For spreadsheet systems to be trustworthy, they often need to be responsive to a multiplicity of criteria that are of value to interested parties. Approaches which enhance spreadsheet trustworthiness can reduce negative spreadsheet risks. This Framework articulates the following **characteristics of trustworthy spreadsheets** and offers guidance for addressing them. Characteristics of trustworthy spreadsheet systems include: **valid and reliable, safe, secure and resilient, accountable and transparent, explainable and interpretable, privacy-enhanced, and fair with harmful bias managed.**

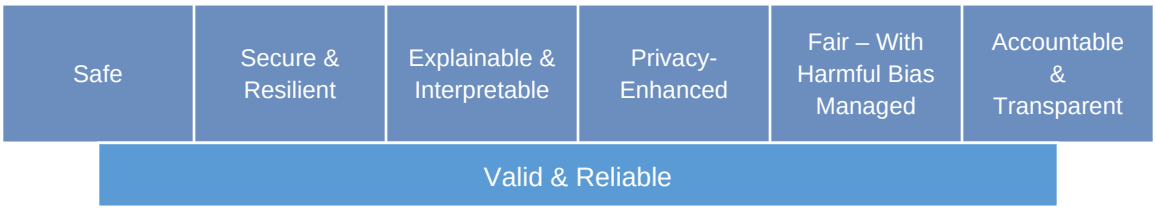

Fig. 2. Characteristics of trustworthy spreadsheet systems. Valid & Reliable is a necessary condition of trustworthiness and is shown as the base. Accountable & Transparent relates to all other characteristics.

Creating trustworthy spreadsheets requires balancing each of these characteristics based on the spreadsheet system's context of use. While all characteristics are socio-technical system attributes, accountability and transparency also relate to the processes and activities internal to a spreadsheet system and its external setting. Neglecting these characteristics can increase the probability and magnitude of negative consequences.

3.1 Valid and Reliable

Validation is the 'confirmation, through the provision of objective evidence, that the requirements for a specific intended use or application have been fulfilled.' Deployment of spreadsheet systems which are inaccurate, unreliable, or poorly tested creates and increases negative spreadsheet risks and reduces trustworthiness.

Reliability is the 'ability of an item to perform as required, without failure, for a given time interval, under given conditions.' Reliability is a goal for overall correctness of spreadsheet operation under the conditions of expected use. This includes the condition where someone inserts a column, breaking all VLOOKUP references to the right of the insertion point. This also includes the condition where the spreadsheet is opened on a machine with different regional settings, causing all decimal separators to be misinterpreted.

Validity and reliability for deployed spreadsheet systems are often assessed by ongoing testing or monitoring that confirms a system is performing as intended. In practice, such testing rarely occurs. The most common validation strategy remains 'it produces a number, and that number seems plausible.'

3.2 Safe

Spreadsheet systems should 'not under defined conditions, lead to a state in which human life, health, property, or the environment is endangered.' Safe operation of spreadsheet systems is improved through responsible design, development, and deployment practices; clear information to deployers on responsible use of the system; responsible decision-making by deployers and end users; and explanations and documentation of risks based on empirical evidence of incidents.

Different types of safety risks may require tailored spreadsheet risk management approaches. Safety risks that pose a potential risk of serious harm – such as spreadsheets used for contact tracing during a pandemic or for determining patient medication dosages – call for the most urgent prioritization and most thorough risk management process. Excel's row limit of 1,048,576 should be considered a safety constraint when processing

datasets that may approach this limit.

3.3 Secure and Resilient

Spreadsheet systems may be said to be resilient if they can withstand unexpected adverse events or unexpected changes in their environment or use. Common security concerns relate to unauthorized access, formula injection, and the exfiltration of sensitive data through spreadsheet sharing. Spreadsheet systems that can maintain confidentiality, integrity, and availability through protection mechanisms may be said to be secure.

Security and resilience are related but distinct characteristics. Resilience relates to robustness and includes the ability to function correctly when a user accidentally deletes a row, when a linked external data source becomes unavailable, or when the spreadsheet is opened in a different version of Excel. The password protection features built into most spreadsheet software should not be considered robust security controls.

3.4 Accountable and Transparent

Trustworthy spreadsheets depend upon accountability. Accountability presupposes transparency. Transparency reflects the extent to which information about a spreadsheet system and its outputs is available to individuals interacting with such a system. Meaningful transparency provides access to appropriate levels of information based on the stage of the spreadsheet lifecycle and tailored to the role or knowledge of spreadsheet actors.

This characteristic's scope spans from design decisions and input data to formula logic, the structure of the workbook, its intended use cases, and how deployment decisions were made and by whom. Transparency is often necessary for actionable redress related to spreadsheet outputs that are incorrect. A transparent system is not necessarily an accurate system; however, it is difficult to determine whether an opaque system possesses such characteristics. Hidden columns, locked cells without visible formulas, and the general tendency to 'just trust the spreadsheet' all undermine transparency.

3.5 Explainable and Interpretable

Explainability refers to a representation of the mechanisms underlying spreadsheet systems' operation, whereas interpretability refers to the meaning of spreadsheet systems' output in the context of their designed functional purposes. Together, explainability and interpretability assist those operating or overseeing a spreadsheet system to gain deeper insights into the functionality and trustworthiness of the system.

Risk from lack of explainability may be managed by describing how spreadsheet formulas function, with descriptions tailored to the user's role and knowledge level. Formula complexity metrics should be tracked:

- **Nesting depth:** How many levels of IF statements are present?
- **Reference sprawl:** How many different sheets are referenced?
- **Volatility:** Does the formula use NOW(), RAND(), or OFFSET()?
- **Opacity score:** How long would it take a new person to understand this formula?

3.6 Privacy-Enhanced

Privacy refers generally to the norms and practices that help to safeguard human autonomy, identity, and dignity. Privacy values such as anonymity, confidentiality, and control generally should guide choices for spreadsheet system design, development, and deployment.

Privacy-enhancing technologies ('PETs') for spreadsheets have been proposed, including approaches adapted from more sophisticated data systems. These include:

Homomorphic Encryption: A theoretical approach that would allow computations on encrypted spreadsheet data without decrypting it. In practice, this technology remains computationally prohibitive for real-world spreadsheet use cases and is primarily discussed in academic contexts. No mainstream spreadsheet software supports homomorphic encryption. Organizations that claim to be 'exploring' this approach are typically not implementing it.

Differential Privacy: A technique that adds calibrated noise to data or query results to prevent identification of individuals. When applied to spreadsheets, this means deliberately making your data wrong in controlled ways. The mathematical guarantees are sound; the practical utility for most spreadsheet use cases is questionable. An organization that implements differential privacy in a budget spreadsheet will have a budget that is, by design, incorrect.

Data Minimization: The practice of collecting and retaining only the data necessary for the intended purpose. This is the only privacy-enhancing approach that is both practical and effective for spreadsheet systems. Unfortunately, it conflicts with the common practice of copying entire datasets 'just in case' and the tendency to never delete any data because 'we might need it later.'

> Under certain conditions, privacy-enhancing techniques can result in a loss in accuracy, affecting decisions about fairness and other values. This is not a bug; it is a fundamental tradeoff. When someone proposes using PETs for spreadsheet data, it is appropriate to ask: 'How much accuracy are we willing to sacrifice?' and 'What decisions will be made with this intentionally-degraded data?' The answer is often that the organization is not actually willing to sacrifice accuracy, at which point the PET discussion becomes purely theoretical.

3.7 Fair – with Harmful Bias Managed

Fairness in spreadsheets includes concerns for equality and equity by addressing issues such as harmful bias and discrimination. NIST has identified categories of spreadsheet bias to be considered and managed:

Data bias: The data used in spreadsheets may not represent the full population or may reflect historical inequities. A spreadsheet model trained on historical hiring data will encode historical hiring biases.

Formula bias: The choices made in constructing formulas – which variables to include, how to weight them, what thresholds to set – all encode human judgments and can perpetuate or amplify bias.

Selection bias: The tendency to exclude certain rows, filter certain conditions, or cherry-pick data ranges can introduce systematic bias into spreadsheet outputs.

Part 2: Core and Profiles

5. SRMF Core

The SRMF Core provides outcomes and actions that enable dialogue, understanding, and activities to manage spreadsheet risks and responsibly develop trustworthy spreadsheet systems. The Core is composed of four functions: GOVERN, MAP, MEASURE, and MANAGE. Each of these high-level functions is broken down into categories and subcategories.

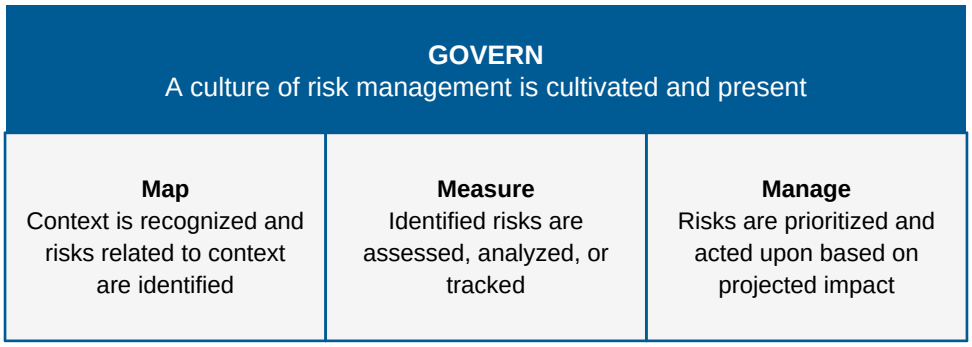

Fig. 3. Functions organize spreadsheet risk management activities at their highest level. Governance is designed to be a cross-cutting function to inform and be infused throughout the other three functions.

5.1 Govern

The GOVERN function cultivates and implements a culture of risk management within organizations designing, developing, deploying, evaluating, or acquiring spreadsheet systems. GOVERN is a cross-cutting function that is infused throughout spreadsheet risk management and enables the other functions of the process.

Strong governance can drive and enhance internal practices and norms. Governing authorities can determine the overarching policies that direct an organization's mission, goals, values, culture, and risk tolerance. Senior leadership sets the tone for risk management within an organization. Management aligns the technical aspects of spreadsheet risk management to policies and operations. Documentation can enhance transparency, improve human review processes, and bolster accountability.

Categories	Subcategories
GOVERN 1: Policies, processes, procedures, and practices across the organization related to the mapping, measuring, and managing of spreadsheet risks are in place, transparent, and implemented effectively.	GOVERN 1.1: Legal and regulatory requirements involving spreadsheets are understood, managed, and documented. GOVERN 1.2: The characteristics of trustworthy spreadsheets are integrated into organizational policies. GOVERN 1.3: Processes are in place to determine needed risk management activities based on risk tolerance. GOVERN 1.4: The organization has established naming conventions that are documented and sporadically followed. GOVERN 1.5: Mechanisms are in place to inventory spreadsheet systems. Most organizations skip this step. GOVERN 1.6: Processes are in place for decommissioning spreadsheets safely. This includes asking 'is anyone still using this?' before deleting it.

GOVERN 2: Accountability structures are in place so that appropriate teams and individuals are empowered, responsible, and trained for managing spreadsheet risks.	GOVERN 2.1: Roles and responsibilities related to spreadsheet risks are documented and clear to individuals throughout the organization. GOVERN 2.2: Personnel receive spreadsheet risk management training. Training should include the knowledge that Excel's date system incorrectly assumes 1900 was a leap year. GOVERN 2.3: Executive leadership takes responsibility for decisions about risks. This is typically assigned to whoever made the mistake of demonstrating competence with Excel.
GOVERN 3: Workforce diversity, equity, inclusion, and accessibility processes are prioritized.	GOVERN 3.1: Decision-making is informed by a diverse team including those who will question whether this should be a spreadsheet at all. GOVERN 3.2: Policies define and differentiate roles for human oversight of spreadsheet systems.
GOVERN 4: Organizational teams are committed to a culture that considers and communicates spreadsheet risk.	GOVERN 4.1: Organizational policies foster a critical thinking and safety-first mindset. GOVERN 4.2: Teams document the risks and potential impacts of spreadsheets they design and deploy. GOVERN 4.3: Practices enable spreadsheet testing, identification of incidents, and information sharing. The instinct to quietly fix errors and hope no one notices is understandable but inadvisable.
GOVERN 5: Processes are in place for robust engagement with relevant spreadsheet actors.	GOVERN 5.1: Policies collect, consider, and integrate feedback from those external to the team that developed the spreadsheet. GOVERN 5.2: Mechanisms enable regular incorporation of feedback. 'The spreadsheet says no' should not be the end of an appeal process.
GOVERN 6: Policies address spreadsheet risks arising from third-party software, data, and supply chain issues.	GOVERN 6.1: Policies address risks associated with third-party data or linked workbooks that may or may not still exist. GOVERN 6.2: Contingency processes handle failures in third-party data sources, including that one cell that references a file on Dave's laptop.

Table 1: Categories and subcategories for the GOVERN function.

5.2 Map

The MAP function establishes the context to frame risks related to a spreadsheet system. The information gathered while carrying out the MAP function enables negative risk prevention and informs decisions for processes such as spreadsheet management, as well as an initial decision about appropriateness or the need for a spreadsheet solution at all.

After completing the MAP function, Framework users should have sufficient contextual knowledge about spreadsheet system impacts to inform an initial go/no-go decision about whether to design, develop, or deploy a spreadsheet system. This includes the fundamental question: should this be a spreadsheet at all?

Categories	Subcategories
MAP 1: Context is established and understood.	MAP 1.1: Intended purposes, potentially beneficial uses, and prospective settings are understood and documented. This includes documenting in a location that is not the spreadsheet itself. MAP 1.2: The spreadsheet's relationship to organizational objectives is understood, even if that relationship is 'someone requested it once.' MAP 1.3: Technical dependencies are identified, including external data sources and linked workbooks. MAP 1.4: The business value has been clearly defined or re-evaluated. MAP 1.5: Organizational risk tolerances are determined and documented.
MAP 2: Categorization of the spreadsheet system is performed.	MAP 2.1: The specific tasks the spreadsheet will support are defined (e.g., budgeting, modeling, tracking). MAP 2.2: Information about the spreadsheet's knowledge limits is documented. If your spreadsheet approaches 1,048,576 rows, you have made decisions that warrant examination. MAP 2.3: Complexity tier is determined: • **Tier 1 (Low):** Single sheet, no external references, affects <10 people • **Tier 2 (Medium):** Multiple sheets, some lookups, affects 10-1,000 people • **Tier 3 (High):** External data connections, macros, affects >1,000 people • **Tier 4 (Critical):** Any spreadsheet used to determine economic policy for nations
MAP 3: Spreadsheet capabilities, targeted usage, and expected benefits and costs are understood.	MAP 3.1: Potential benefits of intended spreadsheet functionality are examined and documented. MAP 3.2: Potential costs, including non-monetary costs from errors or failures, are examined. It is easier to undo a formula than to un-implement an austerity policy. MAP 3.3: Targeted application scope is specified based on the system's capability. MAP 3.4: Processes for practitioner proficiency are defined and documented.
MAP 4: Risks and benefits are mapped for all components including third-party data.	MAP 4.1: Approaches for mapping risks of components – including third-party data – are documented. MAP 4.2: Internal risk controls for components are identified.

MAP 5: Impacts to individuals, groups, communities, organizations, and society are characterized.	MAP 5.1: Likelihood and magnitude of each identified impact are documented. See the Spreadsheet Incident Database for historical examples. MAP 5.2: Power asymmetries between spreadsheet developers and those affected by outputs are recognized. The person who made the model is rarely the person whose benefits are determined by it.

Table 2: Categories and subcategories for the MAP function.

5.3 Measure

The MEASURE function employs quantitative, qualitative, or mixed-method tools, techniques, and methodologies to analyze, assess, benchmark, and monitor spreadsheet risk and related impacts. Spreadsheet systems should be tested before their deployment and regularly while in operation. In practice, the most common validation strategy remains 'it produces a number, and that number seems plausible.'

Categories	Subcategories
MEASURE 1: Appropriate methods and metrics are identified and applied.	MEASURE 1.1: Approaches and metrics for measurement of spreadsheet risks are selected starting with the most significant risks. MEASURE 1.2: Appropriateness of metrics and effectiveness of existing controls are regularly assessed. MEASURE 1.3: Internal experts who did not create the spreadsheet are involved in assessments. The author cannot effectively review their own work; their brain will see what they intended to write, not what they actually wrote.
MEASURE 2: Spreadsheet systems are evaluated for trustworthy characteristics.	MEASURE 2.1: Test sets, metrics, and tools used during testing are documented. MEASURE 2.2: The spreadsheet is demonstrated to be valid and reliable. 'It works on my machine' is not a validation strategy. MEASURE 2.3: Edge cases are explicitly tested including: empty inputs, zero values, negative numbers, values that could be interpreted as dates, values that could be interpreted as scientific notation, and the letter 'E' followed by any number. MEASURE 2.4: The spreadsheet is evaluated for safety risks including row limit breaches. MEASURE 2.5: Security and resilience are evaluated including password protection limitations. MEASURE 2.6: Formula complexity is assessed: nesting depth, reference sprawl, volatility, opacity score. MEASURE 2.7: Privacy risk is examined including inappropriate data retention. MEASURE 2.8: Fairness and bias are evaluated. Remember that SEPT1 is a gene name, not September 1st.
MEASURE 3: Mechanisms for tracking identified spreadsheet risks over time are in place.	MEASURE 3.1: Approaches and documentation are in place to regularly identify and track existing and emergent risks. MEASURE 3.2: Risk tracking approaches are considered for settings where risks are difficult to assess. MEASURE 3.3: Feedback processes for end users to report problems are established. These should not simply direct users to 'check the formula.'

MEASURE 4: Feedback about efficacy of measurement is gathered and assessed.	MEASURE 4.1: Measurement approaches are connected to deployment contexts and informed by domain experts. MEASURE 4.2: Measurement results are informed by input from domain experts to validate whether the system is performing as intended. MEASURE 4.3: Performance improvements or declines based on field data are identified and documented.

Table 3: Categories and subcategories for the MEASURE function.

5.4 Manage

The MANAGE function entails allocating risk resources to mapped and measured risks on a regular basis. Risk treatment comprises plans to respond to, recover from, and communicate about incidents or events. It is incumbent on Framework users to continue applying the MANAGE function as methods, contexts, risks, and needs evolve over time.

Categories	Subcategories
MANAGE 1: Spreadsheet risks based on assessments from the MAP and MEASURE functions are prioritized, responded to, and managed.	MANAGE 1.1: A determination is made as to whether the spreadsheet achieves its intended purposes and whether deployment should proceed. MANAGE 1.2: Treatment of documented risks is prioritized based on impact, likelihood, and available resources. MANAGE 1.3: Responses to high priority risks are developed and documented. Response options include: • **Accept:** Acknowledge the risk and proceed (document this decision) • **Mitigate:** Implement controls to reduce risk • **Transfer:** Have someone else be responsible for the spreadsheet • **Avoid:** Recognize that this should not be a spreadsheet MANAGE 1.4: Negative residual risks to end users are documented.
MANAGE 2: Strategies to maximize spreadsheet benefits and minimize negative impacts are planned and implemented.	MANAGE 2.1: Resources required to manage risks are taken into account – along with viable non-spreadsheet alternatives. MANAGE 2.2: Mechanisms are in place to sustain the value of deployed spreadsheets. MANAGE 2.3: Procedures respond to and recover from previously unknown risks when identified. MANAGE 2.4: Mechanisms are in place to supersede, disengage, or deactivate spreadsheets that demonstrate outcomes inconsistent with intended use.
MANAGE 3: Spreadsheet risks and benefits from third-party entities are managed.	MANAGE 3.1: Risks and benefits from third-party resources are regularly monitored. MANAGE 3.2: Linked workbooks and external data sources are monitored as part of regular maintenance.
MANAGE 4: Risk treatments, including response and recovery, and communication plans are documented and monitored.	MANAGE 4.1: Post-deployment monitoring plans are implemented, including mechanisms for capturing user input, incident response, recovery, and change management. MANAGE 4.2: Activities for continual improvements are integrated into updates. MANAGE 4.3: Incidents and errors are communicated to relevant actors including affected communities. The instinct to quietly fix the error and hope no one notices is understandable but inadvisable. MANAGE 4.4: The organization maintains institutional memory of spreadsheet incidents, recognizing that those who do not learn from VLOOKUP errors are condemned to repeat them.

Table 4: Categories and subcategories for the MANAGE function.

6. SRMF Profiles

SRMF use-case profiles are implementations of the SRMF functions, categories, and subcategories for a specific setting or application based on the requirements, risk tolerance, and resources of the Framework user. Profiles may illustrate how risk can be managed at various stages of the spreadsheet lifecycle.

Minimum Viable Governance Profile

For organizations beginning their spreadsheet risk management journey:

- Someone knows where the important spreadsheets are
- Critical spreadsheets have identifiable owners
- File names provide some indication of content
- The most important spreadsheets have been backed up at least once
- At least one person has asked 'should this be a spreadsheet?'

Responsible Practice Profile

For organizations seeking comprehensive spreadsheet governance:

- All Tier 3+ spreadsheets are documented
- Formula review is conducted before deployment
- Testing occurs with known test cases
- Version control is maintained
- Ownership and succession are defined
- Incident response procedures exist
- Periodic audits are conducted
- Someone has genuinely considered whether the thing should be a spreadsheet

Appendix B: How Spreadsheet Risks Differ from Traditional Software Risks

As with traditional software, risks from spreadsheet-based technology can span organizations and lead to societal impacts. Spreadsheet systems also bring a set of risks that are not comprehensively addressed by current risk frameworks. Compared to traditional software, spreadsheet-specific risks that are new or increased include the following:

- The data used for spreadsheet calculations may not be a true or appropriate representation of the context or intended use.
- Intentional or unintentional changes may fundamentally alter spreadsheet system behavior without triggering any visible error.
- Spreadsheet scale and complexity (thousands of cells, any of which may contain formulas) distributed across a tabular structure rather than visible code.
- Higher degree of difficulty in predicting failure modes due to silent errors.
- Privacy risk due to tendency to copy and retain entire datasets 'just in case.'
- Spreadsheet systems may require more frequent maintenance due to data, formula, or context drift.
- Increased opacity due to hidden columns, locked cells, and distributed formula logic.
- Underdeveloped testing standards and inability to document spreadsheet-based practices to the standard expected of traditionally engineered software.
- Difficulty in performing regular testing, or determining what to test, since spreadsheets are not subject to the same controls as traditional code development.
- Inability to predict or detect the side effects of spreadsheet systems beyond visual inspection of outputs.
- Version proliferation: 'Budget_FINAL_v3_REAL_USE_THIS_ONE(2).xlsx' represents a failure mode unique to spreadsheet systems.

Appendix C: Privacy Enhancing Technologies for Spreadsheets

Privacy Enhancing Technologies (PETs) have been proposed as solutions for protecting sensitive data in spreadsheet systems. This appendix provides a candid assessment of their practical applicability.

C.1 Homomorphic Encryption

Homomorphic encryption allows computations to be performed on encrypted data without decrypting it. In theory, this would allow spreadsheet calculations to be performed on encrypted values, with only the final result decrypted by authorized parties.

Current State:

- No mainstream spreadsheet software supports homomorphic encryption.
- Computational overhead makes real-time spreadsheet use impractical. A simple SUM operation that normally takes milliseconds would take seconds or minutes.
- Only certain operations are supported; complex spreadsheet functions cannot be easily translated.
- Key management complexity adds significant operational burden.
- Organizations that claim to be 'exploring' homomorphic encryption for spreadsheets are typically not implementing it.

Recommendation: Homomorphic encryption is not currently a practical solution for spreadsheet privacy. Organizations should focus on access controls, data minimization, and appropriate handling procedures instead.

C.2 Differential Privacy

Differential privacy adds calibrated statistical noise to data or query results to prevent identification of individuals while preserving aggregate statistical properties.

Practical Implications for Spreadsheets:

- Differential privacy means deliberately making your data wrong in controlled ways. The mathematical guarantees are sound; the practical utility for most spreadsheet use cases is questionable.
- An organization that implements differential privacy in a budget spreadsheet will have a budget that is, by design, incorrect.
- The privacy-accuracy tradeoff must be explicitly acknowledged. How much error is acceptable?
- Repeated queries on the same data consume 'privacy budget,' limiting utility over time.
- Most spreadsheet users lack the statistical expertise to properly implement and interpret differentially private outputs.

Recommendation: Differential privacy may be appropriate for specific use cases involving aggregate statistics on sensitive populations. It is not appropriate for spreadsheets where accuracy of individual values matters. When someone proposes using differential privacy for spreadsheet data, ask: 'How much accuracy are we willing to sacrifice?' and 'What decisions will be made with this intentionally-degraded data?'

C.3 Secure Multi-Party Computation

Secure multi-party computation (MPC) allows multiple parties to jointly compute a function over their inputs while keeping those inputs private.

Practical Assessment:

- MPC requires coordination among multiple parties, which contradicts the typical single-user spreadsheet workflow.
- Communication overhead is significant; real-time collaboration would be impractical.
- Implementation complexity is far beyond typical spreadsheet user capabilities.
- No mainstream spreadsheet software supports MPC protocols.

C.4 What Actually Works

For practical spreadsheet privacy, the following approaches are recommended:

- **Data Minimization:** Collect and retain only the data necessary for the intended purpose. This is the only privacy-enhancing approach that is both practical and effective. Unfortunately, it conflicts with the common practice of copying entire datasets 'just in case.'
- **Access Controls:** Limit who can view or edit the spreadsheet. This is simple and effective.
- **Aggregation:** Use aggregate statistics rather than individual-level data where possible.
- **Anonymization:** Remove or mask identifying information. Be aware that anonymization is harder than it appears; combinations of quasi-identifiers can re-identify individuals.
- **Appropriate Handling:** Don't email spreadsheets containing sensitive data. Don't store them on personal devices. Don't leave them open on your screen during meetings.

> The most effective privacy-enhancing technology for spreadsheets is the decision not to collect sensitive data in the first place. The second most effective is the decision not to put sensitive data in a spreadsheet. Everything else is a distant third.

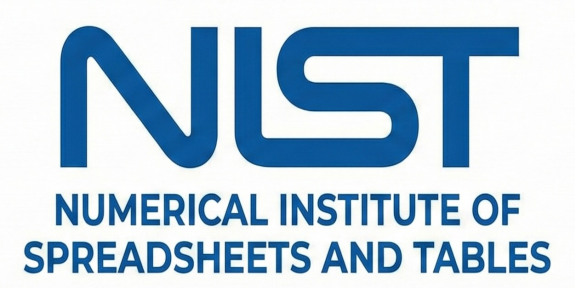

This publication is available free of charge from:
https://doi.org/10.6028/NIST.SS.100-1
(This is not a real URL)

"Because every VLOOKUP is a policy decision."