


TRANSLATIONAL ARTICLE

Using the Five Safes to structure economic evaluations of data governance

Felix Ritchie  and Damian Whittard

Department of Accounting, Economics and Finance, Bristol Business School, University of the West of England, Bristol, UK

Corresponding author: Felix Ritchie; Email: felix.ritchie@uwe.ac.uk

Received: 20 October 2022; **Revised:** 07 February 2024; **Accepted:** 14 February 2024

Keywords: confidentiality; cost-benefit analysis; data governance; economic evaluation; Five Safes

Abstract

As the world has become more digitally dependent, questions of data governance, such as ethics, institutional arrangements, and statistical protection measures, have increased in significance. Understanding the economic contribution of investments in data sharing and data governance is highly problematic: outputs and outcomes are often widely dispersed and hard to measure, and the value of those investments is very context-dependent. The “Five Safes” is a popular data governance framework. It is used to design and critique data management strategies across the world and has also been used as a performance framework to measure the effectiveness of data access operations. We report on a novel application of the Five Safes framework to structure the economic evaluation of data governance. The Five Safes was designed to allow structured investigation into data governance. Combining this with more traditional logic models can provide an evaluation methodology that is practical, reproducible, and comparable. We illustrate this by considering the application of the combined logic model-Five Safes framework to data governance for agronomy investments in Ethiopia. We demonstrate how the Five Safes was used to generate the necessary context for a more traditional quantitative study, and consider lessons learned for the wider evaluation of data and data governance investments.

Policy Significance Statement

Investments in data are seen by governments as central to productivity. Investments in the governance of that data have received less attention, but the pandemic showed why good governance procedures matter. Public sector organizations rely heavily on economic evaluation to infer the value of such investments. Attempts to do this for “data” are fairly meaningless, and for “data governance” they are nonexistent. We describe an approach for evaluating investments in good data governance, using a case study from Ethiopia. Key is the use of the “Five Safes,” a popular framework used for designing and reviewing data governance but not (so far) for evaluation. We argue that this structure has significant potential for supporting such evaluations in data governance, by applying a well-understood framing mechanism to improve the design of evaluation models.

1. Introduction

Data are increasingly seen as a core component of business, even in businesses that traditionally might not be thought of as data-heavy, such as agriculture, food processing, and construction (Saguy and Karel,

1980; Bilal et al., 2016; Kamilaris et al., 2017). Government strategies see data as a key driver of productivity, and public sector data can be among the most useful.

New data protection regulations and the value of sharing data with others have raised the importance of data governance (the processes and procedures to ensure that data use and sharing is ethical, fair and appropriate). Data governance is often seen as a cost, but there are arguments that it should be seen as an investment (Ritchie, 2021; Green and Ritchie, 2023), and there is clear evidence of the ability of good data governance to reduce costs (e.g., Alves and Ritchie, 2020). The COVID-19 pandemic proved a strong test for data governance systems in both the public and private sectors. While many organizations found their data governance processes unsuited to home working, large research organizations with strong effective data governance systems made the switch relatively easily. For example, some secure research data centers (RDCs) were able to easily adjust from on-site access to home working, a policy that had been unthinkable a year before (South and O'Donnell, 2021). In contrast, similar RDCs in other situations simply closed as their data governance systems were not adaptable to the new situation.

When seeing data governance as an investment, it is fair to ask about the return on that investment, and which investments generate the most overall gains. However, putting a value on data governance is perhaps more challenging than valuing data use: governance itself is not the end product, it combines many linked activities of an organization, and the costs of poor governance may not be appreciated if they lead to higher risks but no change in the realized outcomes (e.g., a data sharing strategy has a dubious legal basis but is not challenged during the life time of the project).

In the public sector, with few market interactions to provide direct estimates of value, economic evaluation (EE) methods are used to make decisions about investment. Estimates of the value to the economy of data sharing and use vary so wildly as to make them almost meaningless (Wdowin and Diepeveen, 2020; Alves et al., 2021), although some effective estimates of the value of data can be made in specific cases (Technopolis, 2020). Estimates of the value of good data *governance* are effectively nonexistent.

When considering the case of low- and middle-income countries (LMICs) the problems multiply. Foreign aid investments (and requirements to assess the value of that aid) take place in environments that may have very different cultural preferences to the high-income countries (HICs) where the theory of EE and the dominant privacy perspectives have been developed.

In 2019, the University of the West of England (UWE) was commissioned to review the value of data governance in Bill and Melinda Gates Foundation (BMGF)-funded projects, using a specific low-income agricultural context as a test case. Part of the project aim was to explore *how* data governance could be valued: can one define this complex concept in a meaningful way that can be easily understood and related to the measurement of success and failure?

One popular way of dealing with the multidimensionality of data governance is the “Five Safes” model (Ritchie, 2017). This breaks down data governance into five separate but linked domains: projects, people, settings, outputs, and data. It is widely used in government data access for planning, describing and implementing data governance solutions, for training, and increasingly for legislation (see Ritchie, 2021 for details of all of these). The Five Safes has even been used to structure reviews of data governance (e.g., Ritchie, 2009, 2020; ONS, 2011; Green and Ritchie, 2016), but these have not included formal economic assessments.

The Five Safes was designed to clarify the concepts involved in data governance, and make them separable. The UWE/BMGF project decided to explore whether the Five Safes concept could provide useful structure to a formal evaluation. The resulting report, Whittard et al. (2021), showed that there were both pros and cons to this process, but overall this seemed to be a useful addition to the canon of evaluation techniques. Specifically, placing an explicitly qualitative and subjective framework on top of the traditional quantitative analysis provided a structure for identifying the key factors and the feasibility of assessment.

This article reports on that experiment, and considers the translational lessons learned; specifically

- Is this a useful way to frame evaluations of data governance?
- Does it bring clarity to the process?
- What practical lessons can be learned from this initial trial?

The complete report can be found at Gates Open Research (Whittard et al., 2021). A summary of the EE, which is not our focus here, is given by Whittard et al. (2022).

2. Literature review

2.1. Economic and social evaluation and its application to data investments and data governance

The role of the evaluation is to identify the benefits and costs of a project or intervention; assess whether the intervention is worthwhile and has delivered its aims and objectives; to communicate these results; and, ultimately, to incorporate into a broader project cycle so that lessons can be learned (Alves et al., 2021). The principle of economic and social evaluation (ESE) is well established; for example, the “Magenta Book” (HMT, 2020) is a long-established guide to evaluators in the United Kingdom.

Although the literature talks of “ESE,” in practice evaluation is almost always just the economics. As Alves et al. (2021) note, EE is hard enough without the additional uncertainties brought about by considering less easily measured social costs and benefits.

Even for EE, this can be difficult. The Green Book (HMT, 2022) suggests that EE can take the form of a cost-effectiveness analysis, multi-criteria analysis, economic impact assessment, or cost benefit analysis; and ideally, the choice is between two alternatives considered *ex ante*. Baker (2000) suggests the minimum is a simple counterfactual: “doing X” versus “not doing X.”

Whittard et al. (2021) summarize the guidelines on EE as follows:

- The full costs should include direct and indirect costs and attributable overheads.
- All benefits, both direct and indirect, should be valued unless it is clearly not practical to do—however, it remains important to consider valuing the differences between the options.
- All costs and benefits should be valued at market price (opportunity costs).
- Where possible, estimates for wider social and environmental costs and benefits, for which there is no market price, should be included.
- The value should cover the useful lifetime of the asset encompassed.
- Cost and benefits should be expressed in “real terms.”
- Cost and benefits should be discounted due to private/social time preference.
- Cost and benefits of different options should be valued and the net cost and benefit calculated.

This is an ideal world; these guidelines reflect the fact that methods for EE were designed for large investment projects with easily identified inputs and outputs, and a clear alternative. However, as Alves et al. (2021) discuss, these rapidly run into problems when consider more intangible investments such as those relating to data.

Multiple authors have highlighted the fundamental measurement problem in valuing the benefits of data investments and use. Key confounders include the complementarity between data, infrastructure and processes; the ability for data to be reused infinitely and simultaneously without reducing the “stock”; and the possibility that data appreciates in value over time, rather than depreciating as is normal with assets. Wdowin and Diepeveen (2020) summarize both the issues and the wide range of resulting estimates arising from different assumptions, but even these are just for the valuation of data investments. As noted in the introduction, data are not the same as data governance, and there appear to be no evaluations of investments in the latter prior to Whittard et al. (2021).

The issue of assumptions is perhaps the thorniest issue in relation to evaluations. For post hoc evaluations (i.e., the investment has been made, and its economic worth is now being assessed), the vested interest of both the funder and the investment recipient is about producing a “large number” (Alves

et al., 2021). Given the uncertainty embedded throughout the quantitative evaluation process, it does not take much in terms of loosening the parameters to turn a potentially “failing” project into one that records a strong positive return. Alves et al. (2021), who interviewed both delivery teams and stakeholders, found that almost all interviewees “saw relatively little real value in the number” (Alves et al., 2021, p. 24).

When considering data governance, assumptions become even more important. Much of data governance is about reducing risk; that is, lowering the probability of something negative which might happen but ideally (if the project is well designed) will not. In other words, if data governance is well done, the outcome of it is unobservable by design.

Whittard et al. (2021) throw an additional problem into the mix: evaluation in LMICs. While there have been a number of evaluations on projects in LMICs (including agriculture), evaluation methods designed for HICs may not be appropriate in places where discount rates, financing methods, outside options, and institutions differ significantly from the models envisaged. Whittard et al. (2021) argue that there is a strong need to tailor the evaluation to the particular circumstances of the project, and that qualitative evaluations may be just as informative as quantitative evaluations, if not more.

2.2. *The Five Safes: current application and uses*

Data governance and data access are a complex issue, involving questions of law, ethics, IT, statistics, contracts, HR, accreditation, and other factors (Green and Ritchie, 2023). The Five Safes splits data governance and access questions into five separable but comprehensive and interdependent “dimensions of control” (Table 1).

The Five Safes was first created in 2003 at the UK Office for National Statistics to describe data management systems, particularly access to confidential data for research. In the United Kingdom, New Zealand, and Australia, it is used to describe the governance arrangements for most research data release arrangements in the health and social sciences, and for wider government data-sharing arrangements. It is an increasingly common framework across the public sector and academia in North America, Europe, and Japan. From a relatively quiet start, growth has been exponential in the last five years; the biggest stimulus to adoption was the need for innovative approaches to data governance in the pandemic.

The Five Safes is used in formal legislation, such as the UK Digital Economy Act, and in regulation. In the United Kingdom, for example, the Office for Statistics Regulation bases its guidance on the Five Safes (OSR, 2018), while key academic funders (Administrative Data Research UK, Health Data Research UK, The Innovation Hub), all require bidders to “address” the Five Safes in their data management plans. In Australia, the Five Safes has transmuted into five “Data Sharing Principles” as the basis for the new Data Access and Transparency Act 2022.

As the Five Safes has become better known in the last 10 years, it has become more common to use it actively for designing data strategies in the public sector (OECD, 2014; Green and Ritchie, 2016; ICON, 2016; OSR, 2018; Cranswick et al., 2019) and private sector (Security Brief, 2019; Arbuckle and El Emam, 2020). Green and Ritchie (2023) note that the scope of Five Safes has moved away from research data access and has now been applied to areas as diverse as managing HR systems or compliance modelling.

In summary, the Five Safes is a familiar data governance tool in HICs. While there have been some debates in recent years about whether the framework needs to be adapted or extended (Green and Ritchie, 2023, review these debates), in general its use as a framework is uncontroversial. As Green and Ritchie (2023) note, the Five Safes does not itself provide specific guidance: it is a framework for organizing questions, rather than a checklist for answers. However, with the emergence of comprehensive implementation guides such as that being developed in Australia (McEachern, 2021), this is likely to change.

One area that is lacking is evaluation, both process (does the system perform as well as it could?) and economic (does the system generate a positive return on investment?). Ritchie (2017) noted that, in theory, the Five Safes provides a handy structure for evaluations; in practice, ONS (2011) appears to be the only example to date that has used the Five Safes in a formal evaluation, and that was a process evaluation for a system that mirrored ONS’ system. There appear to be no examples of using the Five Safes for EE.

Table 1. *The Five Safes (adapted from Green and Ritchie, 2023)*

Element	Typical question	Example of problems being addressed
Safe projects	Is this appropriate use and management of the data?	What is the purpose of the access to data? Is this an ethical and lawful use of the data? Is there a data management plan in place? What happens to the data at the end of the project?
Safe people	How much can I trust the data users to use it appropriately?	Do the users have the necessary technical skills? Do the users need training in handling confidential data?
Safe settings	How much protection does the physical environment afford to the data?	How is data stored? Are there physical restrictions on the users?
Safe outputs	How much risk is there in the outputs of the access breaching confidentiality?	If the aim of access is to produce statistics, is there any residual risk by, for example, highlighting outliers?
Safe data	Is the level of detail in the data appropriate?	Is there sufficient detail to allow the project to go ahead? Is this excessive detail that is not necessary for the project?

A second gap is the application of the Five Safes in LMICs. It has been adopted by some public sector organizations in Mexico, Nepal, and South Africa, and in 2022, the UK National Institute for Health Research commissioned a 4-year program of virtual “summer schools” for LMIC researchers. However, these particular initiatives have not yet led to a wider adoption of the framework.

3. Combining models for evaluating data governance: a case study

3.1. Project genesis

In 2019, the UWE team was commissioned by CAB International (CABI) to carry out an evaluation of data governance on BMGF projects. The initial aim was simply to provide an EE and to demonstrate how this could be done in practice. This was a retrospective EE, carried out on a completed (or nearly completed) project. An alternative to EE would have been a process evaluation (looking at processes from an operational efficiency perspective), but this was not explored as the aim was to see how EE could be generally employed in BMGF projects.

After initial discussions, it became clear that existing EE models were inappropriate due to the issues raised above: unmeasurability of inputs, outputs and outcomes; conceptual problems with what was being measured; and the need to situate the evaluation in the LMIC context. It was agreed that a new approach would be needed, and the project aims changed to include the development of an effective valuation method, as well as the application of that method to a BMGF project as proof-of-concept.

The evaluation was carried out on Supporting Soil Health Interventions in Ethiopia (SSHIE). This is a \$1.5 million project led by Gesellschaft für Internationale Zusammenarbeit (GIZ). The overall objective of the SSHIE project was to provide proof of concept for the impact of leveraged geo-spatial soil, agronomic, and soil health data on transformative agricultural development. Specific aims of SSHIE were to

- increase the quality, efficiency, and reach of government and private sector services;
- improve national and regional resource allocation and policies; and
- increase the quality, availability, and utility of data and evidence.

Table 2. Detailed governance questions for work packages (subset; from Whittard et al., 2021)

Domain	Sub domain	Example questions
Safe projects	Project planning	• What was the approval process?
	Data management plan	• Did you have a data management plan at beginning?
	Approvals process	• Who was responsible for the approval process?
	Public Engagement	• Is explaining your role to the wider public part of your project's objectives?
Safe people	Governance	• What model and structure of data governance and access is used?
	Training	• What, if any, training is provided for different groups (e.g., data collectors/ IT staff and data processors/users)
	Access	• What systems and procedures are in place to ensure data users operate in an appropriate manner?
Safe data	Data quality	• How is the flow of data processed in the project?
	Data need	• How do you limit the level of detail available depending on the type of user and level of need?
	Compliance	• What are the mechanisms through which breaches of data governance procedures are enforced?
Safe settings		• How do you make data available from a safe setting?
Safe outputs	Disclosure risk	• Are clear standards set and adhered to in relation to disclosure control?
	Building capacity	• Is data created specifically for a project retained after the project is completed?

The aim of the evaluation was not to provide a view on the value of the data acquired, but on the governance arrangements; that is, the value of how the data collection was designed, and use regulated. Accordingly, the evaluation focused on the data management arrangements for the SSHiE project and not on the data collected.

The agreed approach was to split the evaluation of the project's data governance into a qualitative component that would then direct and frame data collection for a quantitative component. Whittard et al. (2021) describe the evaluation in detail. Whittard et al. (2022) summarize and comment on the measurement issues associated with the evaluation. The evaluation contained several innovations. One of those concerns us in this article is the decision to use the Five Safes as a framework to structure the qualitative data collection.

3.2. Method

The first task was to identify the activities in SSHiE that delivered "data governance." This was not identified as a separate activity in the SSHiE project planning, which concentrated on the practical delivery of the data collection mechanisms. The initial discussions between the review team and BMGF/GiZ sought to understand what data governance covered, and to find a common language. This is when we began to consider using the Five Safes as the organizing framework. Three of the five work packages were identified as primarily or mostly about improving the governance of data on the project.

The second stage was to identify what elements of data governance were being addressed in that work package. The Five Safes was broken down into additional subdomains, with the idea that each separate project work package with elements of governance would be tested against the subdomains, where relevant.

Table 3. Typical logic model (HMT, 2020)

Term	Definition
Inputs	Resources used to deliver the project
Activities	What is delivered to the recipient
Outputs	What the recipient does with the activity
Intermediate outcomes	The intermediate outcomes of the project produced by the recipient
Impacts	Wider economic and social outcomes

The third element was the link between activities and outcomes. Every evaluation requires a “logic model,” a statement of how inputs, outputs, impacts and outcomes interact. Table 3 shows a typical logic model suggested by the UK government.

Combining these three elements (work packages, Five Safes, logic model) formed the basis for the team to structure its qualitative investigation.

Table 4 shows a partial view on the table generated for each of the work packages. Thus it can be seen that, for example, in work package 1, one of the activities was the Open Data Sensitization Workshop. Inputs and measurable outputs were identified. The only relevant governance domain is “people,” and this was used to identify the target outcomes.

To undertake the evaluation, a three-staged approach was completed (see Whittard et al., 2021 for fuller details). In summary, the evaluation was based on 3 reports provided by BMGF, 5 interviews/focus groups covering 7 key stakeholders, and 12 targeted questionnaires, of which 8 were completed. Due to budgetary and time limitations, it was not feasible to interview individuals from all stakeholder groups or to identify a representative sample. Instead, a convenience/snowballing approach was taken. The individuals interviewed included representatives from the Ethiopian Government, agriculture sector groups, academics, and the voluntary sector. The questionnaires were sent to experts identified by the CABI team as being able to provide detailed insights into the project.

Measurable inputs and output associated with the activity and the outcomes were identified. This was used to generate specific questions for that activity. A simplified version of the full framework was then sent to the interviewees prior to interview, to indicate the topics that would be discussed, but the role of the Five Safes in structuring the questions was not emphasized to interviewees. The interviews were transcribed and analyzed.

As a post hoc evaluation with limited access to participants, Tables 2–4 represent the starting points of the analysis. Some of the questions were not appropriate or irrelevant, some were felt to be unanswerable by interviewees, and some did not seem likely to lead to useful answers. Nevertheless, the act of defining a general framework before moving to specific topics helped the team rapidly narrow the scope of what was feasible.

3.3. Analyzing findings

The interview findings had two functions: as direct evaluation data and to help structure the quantitative study for the more formal EE.

The direct analysis of the interviews was structured around the Five Safes. For each dimension, the team reported interview findings, and identified lessons learned for that governance dimension. These are tabulated below, in summary, with a subset of findings shown for illustration. For the report, the Five Safes was replaced with more meaningful descriptions of the dimensions (left-hand column of Table 5).

Having used the interviews to determine what factors seemed to help or hinder the successful exploitation of data, the qualitative data collection was designed to put values, as far as possible, on these factors. The interview findings were then reused to qualify and provide context for the quantitative findings.

Table 4. Project mapping (partial view; from Whittard et al., 2021); asterisks indicate which ‘safe’ is relevant for this output

Logic model						Five Safes				
Work package	Inputs	Activity	Outputs	Intermediate outcomes	Long-term outcomes	Project	People	Setting	Outputs	Data
WP:1 Develop Policy	Number of people	Open Data Sensitization Workshop (awareness raising)	Number and type of attendees Policy documents/ guidance produced	Common basic understanding of data access issues among relevant decision-makers so that decision-making is improved	Capacity: individuals who understand value of open data		*			
	Total cost	Development of Coalition of willing	Identification of partners Identification of objectives Identification of responsibilities	Agreement on project aims, providers, beneficiaries, operating methods, so that future discussions do not need to revisit basic points	Example of proactive governance to wider community	*	*			
	Other tangible inputs	Consultation workshop on data access and sharing	Number and type of attendees Policy documents/ guidance produced	Common basic understanding of data access issues among relevant decision-makers reducing time to get agreement in future	Guidelines for the wider community	*	*	*	*	*
	Other intangible inputs	Establishment of data sharing taskforce	Taskforce membership, longevity Policy documents/ guides produced	Recognized authority for project decisions reducing approval time/ activity	Recognized authority for decisions on data sharing in the wider community	*			*	*

Table 5. *Lessons learned (matrix; extract based on Whittard and Ritchie, 2021)*

Lessons learned in respect of...	
...organization of projects and investment in setting up processes	<ul style="list-style-type: none"> • Effective working links between local and international partners was key to maximizing project value • Engagement with a diverse set of stakeholders, even though it may be costly and time consuming, also provides substantial benefits in achieved full project value
...organization of and investment in people	<ul style="list-style-type: none"> • Identify top-level “champions,” and allow time for the relationships to develop • Recruit/empower stakeholders at all levels of delivery to practically drive the project forward
...the set up and management of safe settings	<ul style="list-style-type: none"> • Individuals need to trust the processes for handling data (which could include trust in automated systems) • More attention should be paid to second-order conditions (governance, staff training) when planning IT solutions
...producing statistical outputs and shareable datasets	<ul style="list-style-type: none"> • Outputs produce direct tangible evidence of value created • Training in producing safe outputs is not considered (this is a very common gap in training: Derrick et al., 2022)
...investments in data quality and usability	<ul style="list-style-type: none"> • Working through trusted local organizations and individuals can realize value of the data • Standardization (or, at least, interoperability) of data collection and processing should be prioritized

For example, the quantitative study showed that, when restricting the evaluation to direct, reasonably measurable components, the investments in data governance had a negative benefit (in other words, BMGF would have been better off choosing to invest the money in something other than data governance). As Alves et al. (2021) note, the ability to generate positive returns through highly subjective assumptions makes “the number” all but meaningless when dealing with highly intangible outcomes. We calculated mostly cash and time costs, generating a negative return; we then present the funders with the non-monetary benefits (identified during the qualitative stage) associated with that notional loss, for them to consider whether the overall outcome was valuable. The evaluation then becomes “given these measurable financial outcomes, do the non-financial outcomes seem an appropriate return on the investment?”

However, the information collected on the qualitative section showed that a large number of intangible, persistent benefits had been generated in data governance, even if most of them were thought to be unmeasurable. For example, a key innovation in the project was the Coalition of the Willing (CoW)—an initial meeting and then subsequent activities designed explicitly to build relationships that could make the project work. In the evaluation, this was accounted for on largely cost-of-time measures, but the qualitative review showed that the effects of the CoW were pervasive across the Five Safes: getting stakeholder buy-in, identifying potential champions and blockers, demonstrating the importance of trust in policy-making and delivery. Moreover, interviewees recognized the long-term value of the CoW for changing attitudes, for introducing new ways of thinking about data governance, and for providing a positive example of collective planning. This highlighted that the timing of the evaluation was crucial to the perspective on whether it brought a net gain or not. Some very limited assumptions were sufficient to show that the project overall was of benefit to the community; for example, assuming that the research papers produced during the project were valued using the time-cost of production. This enables consistency with a more traditional measure-everything analysis, if the funders had wanted it.

4. Reflections

4.1. *The Five Safes as a structural framework*

The review team decided to (a) split the review into quantitative and qualitative elements, and (b) use the Five Safes framework to provide structure for the qualitative component. Decision (a) was undoubtedly the more important conceptual choice, as it provided a practical way of identifying the scope of the quantitative study and a context for interpreting quantitative results. However, decision (b) simplified the qualitative element considerably.

The Five Safes does not directly address the measurement or timing of costs and benefits, or many of the other practical problems of evaluation noted in the literature review; as Green and Ritchie (2023) discuss, the Five Safes is a mechanism for framing questions, but provides no direct answers per se. Nevertheless, using a ready-made framework, familiar to the reviewers and with an intellectual hinterland to support it, made the task of breaking one large problem into many smaller ones much easier.

It is worth noting that the qualitative/quantitative split and the use of the Five Safes was not the first solution considered for evaluating the CABI project. Several alternative conceptual frameworks were explored over some 6 months, often based on the more traditional quantitative or case study evaluation models. All proved unsatisfactory, mostly because they did not address the slipperiness of valuing the concept of data governance. Realizing that the Five Safes could provide the scope definition led to the qualitative/quantitative split, and the development of a practical method.

Would other approaches to the qualitative study also have worked? “Yes, probably” is the fairest answer. However, part of the reason for choosing to use the Five Safes is that it was explicitly designed to deal with the multifaceted nature of data governance, and the language of the Five Safes can be directly related to data governance questions. This helped to address the issues of what, exactly, on the SSHiE project counted as “data governance” activities and outcomes. A particular strength of the Five Safes is that it is broad in its coverage, so it ensures that all important elements are captured, while it is subjective in its application. This empowers evaluators to work within a broad framework, while allowing them the flexibility to focus in on issues they feel are of particular importance within the specific context of the project. In addition, given its familiarity to the data science community, using the Five Safes to structure an evaluation can help to reduce some of the uncertainty in an area, which, by its intangible nature, is plagued by ambiguity.

4.2. *Mapping outcomes to the Five Safes*

It was notable that much of the project value came in the “safe projects” dimension: that is, the element of governance concerned with objective-setting, outcomes, approvals, engagement, and planning. This is unsurprising: given the innovative nature of the project and the data governance activities (such as the CoW), planning and clear project oversight were disproportionately important in determining the success of the data governance plans. Ritchie and Tava (2020) make the same point, but on a theoretical rather than empirical basis. It would perhaps be useful, in future, to see whether follow-on projects place the same emphasis on planning, or whether more technical activities such as staff training become the drivers of success or failure.

This does raise the issue of whether those doing the project planning have the necessary skills to do this effectively. Data governance does tend to have strong learning effects; that is, improvements made in governance persist. A framework such as the Five Safes can help to understand where these gains are made on both current and future projects. For example, training and practical experience (the “people” element) in this project should feed directly through into the “project planning” phase of the next investment.

If, as suggested in the qualitative studies, this project has led to permanent changes in attitude toward data sharing and data governance in Ethiopia, then one would expect future projects to reflect this. For example, “a shared understanding of the importance of trust” becomes an input, rather than an output. This is an important outcome for the donor community and would not have been picked up, had the evaluation stuck to a traditional quantitative approach. It is important for the SSHiE project (as well as us) to reflect

and recognize that this program evaluation using the Five Safes might therefore be an easy way to identify accumulative gains from investment (of course, the same could be said for any other structure which is used consistently across multiple evaluations).

One factor that came out strongly in the interviews was the importance of timing. In some of the “safes” benefits were quickly identified as structural shifts that could have long-term impacts (providing a successful example of stakeholder management; showing how attitudes changed); in others, the benefits were more directly related to the projects (such as the relationship between system design and user training). As well as reinforcing the sensitivity of results to the timing of the evaluation, such findings also show the usefulness of a structure that can highlight the different temporal impacts of different elements of governance.

4.3. The value of a consistent framework

The Five Safes is more commonly used to design data governance activities rather than evaluate them, but the implication is there is the potential to exploit this framework consistently throughout all stages of grantmaking (portfolio strategy, grant design, anticipated risk of achieving measurable impact). It was noted that data governance was not considered as an explicit planning process when setting up SSHiE. This is not uncommon; data governance is often seen as a set of separate tasks, such as ethical review or data management, which is why the Five Safes was created. Using the Five Safes, or some other framework, to develop an explicit data governance strategy would have simplified evaluation, even a retrospective one as in this case.

The Five Safes also showed itself as adaptable to the LMIC situation with relatively little change; again, this is because the Five Safes is a framework, not a checklist, and so the details can be tailored to the local situation while keeping the “big picture” constant. This has been the authors’ (limited) experience of applying the Five Safes in other LMIC contexts where, for example, training has seen local adaptation while still retaining the same overall concepts. This flexibility does offer a basis for developing data governance planning and evaluation models across very different contexts.

The Five Safes began as a tool for designing access to confidential data in research facilities. Its growth in this area has followed the typical S-shaped pattern of a network tool: slow initial growth, quick adoption at the point at which it becomes “the” standard, then slowing growth as it reaches saturation. Outside of this specific use, it is still at the slow-growth stage, but the range of applications is increasing. As noted in the literature review, the use of the Five Safes in evaluation is rare, but then the evaluation of data governance itself is very rare. Providing an evaluation tool that is more widely consistent with the language used in data governance is likely to produce some synergy in both evaluation and planning.

5. Limitations and learning

As this was the first attempt to use the Five Safes in this way, there was considerable learning. For example, there were originally far more questions than those presented in [Table 2](#), and these were not structured below the level of the Five Safes. After discussion with CABI about what could be reasonably expected of the interviewees, the team reduced the number and type of questions and added the subdimensions to allow signposting of questions.

The framework was helpful to the researchers in planning for the interviews, and to the CABI team for seeing how the analysis would be structured. However, feedback from the interviewees made clear that the work package-governance matrix was less useful for them, and could even add confusion. Interviewees wanted a clear set of questions that could be answered, and had no interest in the conceptual framework. The lesson from this is that the Five Safes is more of a tool for the evaluators, and bringing the jargon of the Five Safes into data collection may be counterproductive. For example, identifying which of the “safes” was relevant to which work package (as in [Table 4](#)) was supposed to streamline questions; in practice, interviews were allowed to develop organically, and so the structure was only used to write up the result afterward.

Finally, alternative frameworks and evaluation models do exist, and this case study made no attempt to compare the benefits and drawbacks of different approaches. The aim of the project was to find a method that worked on a problem as intangible as the value of data governance investments in LMICs. As noted above, we spent several months trying to identify a useful structure that was both conceptually meaningful and practically achievable. The Five Safes+logic model+qual/quant approach was the result, but the Five Safes was a relatively late addition. It is quite possible that some alternative structure might have worked as well, but a key factor in choosing this approach was the comparability of the Five Safes with other data governance activities and terminology (as well as its simple structure).

6. Conclusion

The public sector relies heavily on formal EE to understand the costs and benefits of investment decisions. Understanding the economic contribution of investments in data sharing and data governance is problematic: outputs and outcomes are often widely dispersed and hard to measure, and the value of those investments is very context dependent. The evaluation of investments in data governance has been almost completely ignored in both practitioner and academic literature, perhaps because of the challenges involved in conceptualization, definition, and measurement.

The SSHiE case study demonstrated that there was considerable value of undertaking a wider evaluation that integrates both quantitative and qualitative frameworks. Although a number of different evaluation models were tried over the 18 months of the project, the Five Safes became the basic structure of the study. It helped to clarify what data governance was, and how to ask questions about it. The qualitative–quantitative split helped define what evaluation data could reasonably be gathered, and how to interpret it in context. There was a significant amount of learning, but the use of the Five Safes opens up a number of interesting guidelines for future evaluations.

First, by using the Five Safes model, the evaluators were able to understand the wider value of improvements in data governance. A traditional quantitative evaluation would have been primarily focused on tangible investments that potentially were (a) siloed and (b) focused on the data/technology (rather than intangibles from data governance). If the Five Safes (or at least some other qualitative framework) had not been employed, quantitative evaluations which fail to pick up the nuanced effects of such projects can reinforce the loop of bad governance, duplicate data capture, lack of trust, lack of sharing, and so on.

Second, the Five Safes framework benefits from ensuring evaluators take a broad view of all important elements of data governance, while being flexible enough in its application to allow evaluators to focus in on project and contextual specific issues—thus improving the validity of the results. Its applied application in the SSHiE case study has demonstrated its compatibility with more traditional and formal quantitative evaluation techniques. This is also consistent with other work to apply the Five Safes in LMICs.

Third, familiarity with the Five Safes in the data science community should encourage a wider acceptance of the method, allowing for assessment of the reliability of findings between studies. The Five Safes is a simple framework, but one that can be applied in many different situations. EEs are notoriously difficult to compare because they are highly dependent on assumptions and data quality: everything from scope to timing to discount rate to counterfactual is the result of a choice by someone, sometimes for no better reason than that one thing is doable but another is not. The Five Safes is explicitly subjective to encourage decision makers to acknowledge these choices; having a consistent, common framework (albeit one where implementations may differ widely) encourages evaluators to see how the dimensions compare to each other across projects. This case study does not prove that the approach will resolve differences, but we suspect, based on our extensive experience in data governance, that others will also find it helpful. Moreover, in a field dominated by uncertainty and subjectivity, a common framework, even one that is itself subjective, does tend to bring clarity over choices and decisions.

Fourth, value in the long term comes from the option of commissioning evaluations at the start of an investment rather than at the end (Alves et al., 2021). The Five Safes is more often used to design data governance systems and processes. By using it to structure evaluations as well, this may allow project managers to see more clearly how each element of their risk management strategy appears to generate value.

The drawback of the Five Safes framework, that as a tool for structuring questions, it has little detail, means that it does not prevent scope creep or the potential for evaluators to apply inappropriate judgments; any tool can be used well or badly, and Green and Ritchie (2023) acknowledge that this applies to the Five Safes as well. No framework by itself is a panacea. However, both greater familiarity with the model and the potential for integration with formal implementation toolkits (such as the HMT, 2020, 2022) suggests that a broad level of comparability in evaluation may be achievable—no small outcome in a field where intangibility is the defining characteristic. While further studies are needed to firm up these ideas, our experience from applying the Five Safes in multiple environments is that the adoption of the common framework in this new area is likely to bring benefits to both analysts and policymakers.

Data availability statement. The participants of this study did not give written consent for their data to be shared publicly, so due to the sensitive nature of the research (commentaries about individual contributions to the project) supporting data are not available.

Author contribution. Both authors contributed equally to the article. F.R. was the lead on funding acquisition, methodology, and writing and supported D.W. on other areas. D.W. was the lead on the original investigation and project management, and supported F.R. on other areas.

Funding statement. The original project on which this study is based was funded by CAB International. The funder had no role in the decision to prepare or publish the manuscript, or in the writing of it.

Competing interest. The authors declare none.

References

- Alves K and Ritchie F (2020) Runners, repeaters, strangers and aliens: Operationalising efficient output disclosure control. *Statistical Journal of the IAOS* 36(4), 1281–1293.
- Alves K, Tava F, Whittard D, Green E, Beata Kreft M and Ritchie F (2021) *Process and Economic Evaluation of the ODI R&D Programme: Final Report*. London: Open Data Institute.
- Arbuckle L and El Emam K (2020) *Building an Anonymization Pipeline: Creating Safe Data*. Toronto: O'Reilly.
- Baker JL (2000) *Evaluating the Impact of Development Projects on Poverty: A Handbook for Practitioners*. Washington, DC: The World Bank.
- Bilal M, Oyedele LO, Qadir J, Munir K, Ajayi SO, Akinade OO, Owolabi HA, Alaka HA and Pasha M (2016) Big data in the construction industry: A review of present status, opportunities, and future trends. *Advanced Engineering Informatics* 30(3), 500–521.
- Cranswick K, Tumpene S and Stobert S (2019) *Virtual Data Labs – A More Flexible Approach to Access Statistics Canada Microdata*. UNECE/Eurostat Work Session on Statistical Data Confidentiality, The Hague, October.
- Derrick B, Green E, Ritchie F, Smith J and White P (2022) *The Inadvertently Revealing Statistic: A Systemic Gap in Statistical Training?* Presentation to the Scottish Economic Society Conference, April.
- Green E and Ritchie F (2016) *Data Access Project: Final Report*. Australian Department of Social Services, June.
- Green E and Ritchie F (2023) The present and future of the five safes framework. *Journal of Privacy and Confidentiality* 13(2), 1–20. <https://doi.org/10.29012/jpc.831>.
- HMT (2020) *The Magenta Book: Guidance for Evaluation*. London: HM Treasury. Available at <https://www.gov.uk/government/publications/the-magenta-book> (accessed March 2024).
- HMT (2022) *The Green Book: Appraisal and Evaluation in Central Government*. London: HM Treasury. Available at <https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government> (accessed March 2024).
- ICON (2016) *Hellenic Statistical Authority: Mission on Microdata Access; Mission Report*. Luxembourg: Eurostat.
- Kamilaris A, Kartakoullis A and Prenafeta-Boldú FX (2017) A review on the practice of big data analysis in agriculture. *Computers and Electronics in Agriculture* 143, 23–37.
- McEachern S (2021) *CADRE Five Safes Framework: Conceptualisation and Operationalisation of the Five Safes Framework*. CADRE (Co-ordinated Access for Data Research and Environments) Project. <https://doi.org/10.5281/zenodo.5748610>.
- OECD (2014) *OECD Expert Group for International Collaboration on Microdata Access: Final Report*. Organisation for Economic Co-operation and Development, July.
- ONS (2011) *Secure Data Service: Risk Assessment*. Newport: Office for National Statistics.
- OSR (2018) *Joining up Data for Better Statistics*. Office for Statistics Regulation Systematic Review Programme Report, September.
- Ritchie F (2009) UK release practices for official microdata. *Statistical Journal of the IAOS* 26(3, 4), 103–111.
- Ritchie F (2017) *The “Five Safes”: A Framework for Planning, Designing and Evaluating Data Access Solutions*. Data for Policy Conference 2017. <http://doi.org/10.5281/zenodo.897821>.

- Ritchie F** (2020) *Review of Data Governance and Training for Nepal Injury Research Centre*. Kathmandu: Nepal Injury Research Centre.
- Ritchie F** (2021) Microdata access and privacy: What have we learned over twenty years? *Journal of Privacy and Confidentiality* 11 (1), 1–8.
- Ritchie F and Tava F** (2020) 'Five Safes' or 'One Plus Four Safes'? *Musing on Project Purpose*. Bristol Centre for Economics and Finance Blog, 27 July. Available at <https://blogs.uwe.ac.uk/economics-finance/>.
- Saguy I and Karel M** (1980) *Modeling of Quality Deterioration during Food Processing and Storage*. Food Technology (USA).
- Security Brief** (2019) *IXUP Embeds Five Safes Framework in Platform*. Available at <https://securitybrief.com.au/story/ixup-embeds-five-safes-framework-in-platform> (accessed March 2024).
- South B and O'Donnell N** (2021) *Responding to the COVID-19 Pandemic to Enable Research to Continue in the Secure Research Service (SRS)*. Conference Presentation "Covid-19: A Catalyst for Greater Data Collection and Access?", January.
- Technopolis** (2020) *Process Evaluation of the ODI's Innovation Programme: Stage 2 Report*. London: ODI.
- Wdowin J and Diepeveen S** (2020) *The Value of Data*. Cambridge: Bennett Institute for Public Policy.
- Whittard D, Ritchie F, Musker R and Rose M** (2022) Measuring the value of data governance in agricultural investments: A case study. *Experimental Agriculture* 58, e8. <https://doi.org/10.1017/S0014479721000314>.
- Whittard D, Ritchie F and Nyengani J** (2021) Measuring the value of improving data governance and access in gates foundation programme: A case study of the supporting soil health interventions in Ethiopia projects. *Gates Open Research* 5, 1–107. <https://doi.org/10.21955/gatesopenres.1116786.1>.