THE OHIO STATE UNIVERSITY
PROGRAM ON DATA AND GOVERNANCE
A program of the **Moritz College of Law** and
the **Translational Data Analytics Institute**

iapp

RESPONSIBLE AI MANAGEMENT

Evolving Practice, Growing Value

THE OHIO STATE UNIVERSITY
PROGRAM ON DATA AND GOVERNANCE
A program of the **Moritz College of Law** and
the **Translational Data Analytics Institute**

iapp

RESPONSIBLE AI MANAGEMENT:

Evolving Practice, Growing Value

By: Professor Dennis Hirsch, Faculty Director, Program on Data and Governance
Jared Ott, Senior Research Associate (fmr.), Program on Data and Governance
Angie Westover-Munoz, Program Manager, Program on Data and Governance

TABLE OF CONTENTS

TABLE OF FIGURES

ACKNOWLEDGEMENTS

Cisco provided funding for this research under its Ethics in AI RFP, https://research.cisco.com/open-rfps/RFP-21-01. We thank Cisco for this support and particularly thank Cisco's Robert Waitman, Lisa Bobbitt, and Chris Shenefiel (formerly Cisco, now William & Mary) for their unflagging interest and support.

The International Association of Privacy Professionals (IAPP) distributed the survey and commented on the survey instrument and this report. The Future of Privacy Forum (FPF), Information Accountability Foundation (IAF), and Centre for Information Policy Leadership (CIPL) each also distributed the survey. We thank each of these organizations for their important contribution. We are particularly grateful for the input that we received from the IAPP's Caitlin Fennessy, Ashley Casovan, Joe Jones, and Lynsey Burke.

Finally, this study would not have been possible without the ongoing support that The Ohio State University Moritz College of Law and the Translational Data Analytics Institute provide for the Program on Data and Governance and its work on the governance of AI.

EXECUTIVE SUMMARY

Faced with growing public and legal pressure, some businesses are taking steps to utilize AI in a more socially responsible way. They refer to these efforts as "responsible AI management" (RAIM). This report conveys the results of a survey-based study, conducted in early 2023, of RAIM practices at businesses that develop and use AI. The study sought to answer three questions (shown below) with the report's main conclusions for each question beneath.

What do businesses do when they pursue RAIM?

Businesses RAIM programs consist of at least 14 different activities, some of which are more common than others. Most RAIM programs are immature and underdeveloped at this point.

Who in the business is responsible for RAIM?

Privacy experts are most likely to be responsible for RAIM. Data analytics, ethics, and risk experts are involved as well.

What value does RAIM create for the companies that implement it?

RAIM initiatives create significant business value. This value comes more from increased product quality and trust, and decreased risk, than from the fulfillment of corporate values.

Eight key findings further explain and support the above insights and suggest areas for additional research.

01 Large companies provided most of the responses

Most respondents came from large companies (1,000 or more employees, $10,000,000 or more in annual revenue). This could mean that larger companies are doing more with respect to RAIM and so were more likely to respond to the survey. If true, this could create selection bias in favor of companies that have more developed RAIM programs. Respondents came from nine different industry sectors. This suggests that the use of AI, and the resulting need for AI governance, is present in many sectors.

02 RAIM includes at least 14 actions

The most commonly reported RAIM activities involved risk assessment (evaluating regulatory risk, identifying harms to stakeholders), the building of a RAIM management structure (responsible official, RAIM committee), and the adoption of substantive RAIM standards (AI ethics principles, RAIM policies). Respondents were much less likely to train employees in RAIM, and to evaluate their employees' or their organization's RAIM performance.

03 Privacy experts are most likely to be responsible for RAIM, with others involved as well

Sixty percent of respondents said that their organization had assigned the RAIM function to a specific person or persons. The people performing this function held a variety of titles ranging from privacy manager, to data scientist, to responsible AI officer. Companies were most likely to assign the RAIM function to individuals with expertise in privacy (59.5%). The number of companies that identified more than one person involved in RAIM, and the wide variety of titles those individuals hold, suggests that a cross-functional approach to RAIM may be useful.

04 Businesses believe that RAIM is important to their company

A majority of the survey respondents (68%) said that RAIM was either important or extremely important to their company. Nearly 90 percent either "Agreed" or "Strongly Agreed" with the idea that companies should make a "meaningful investment" in RAIM.

05 RAIM creates substantial value for companies that invest in it

All respondents reported that their company gets at least some value from their RAIM programs, with almost 40 percent reporting "a lot" or a "great deal" of value. Companies with more developed RAIM programs report gaining greater value, on average.

06 RAIM provides strategic value

RAIM serves a strategic function, in addition to allowing a company to better attain its corporate values. It improves product quality, trust, employee relations, and reduces risk. The data suggest that RAIM may produce the most value by improving product quality. It achieves this by promoting AI product innovation and better meeting customer expectations, not just by reducing negative impacts.

07 Specific RAIM activities may produce particular types of business value

The survey data begins to tease out relationships between particular RAIM activities and specific types of value created. The data suggest that companies that adopt RAIM policies experience greater increases in trust, and those that adopt RAIM policies and require their suppliers to follow them report both increased trust and competitive advantage. In addition, businesses that attempt to identify whether their AI products and processes may cause harm to others or violate the law report greater increases in product quality. The small number of respondents precludes any definitive conclusions about these relationships. Additional research will be required to substantiate and explain them.

08 Implementation lags enthusiasm

Most respondents said that their RAIM programs were at an early stage of implementation, and a majority described their company's process for making responsible AI judgment calls as "ad hoc" rather than "systematic." This halting implementation stands in sharp contrast to respondents' belief in RAIM's importance and perceived value. This contrast signals that further RAIM development may offer an opportunity to capture this value.

INTRODUCTION

People are increasingly coming to see artificial intelligence as a double-edged sword. On the one hand it promises to enhance individual and business performance, cure diseases, solve environmental problems, and otherwise benefit humanity. On the other, it threatens to perpetuate bias, invade privacy, spread misinformation and, according to some, threaten humanity itself.

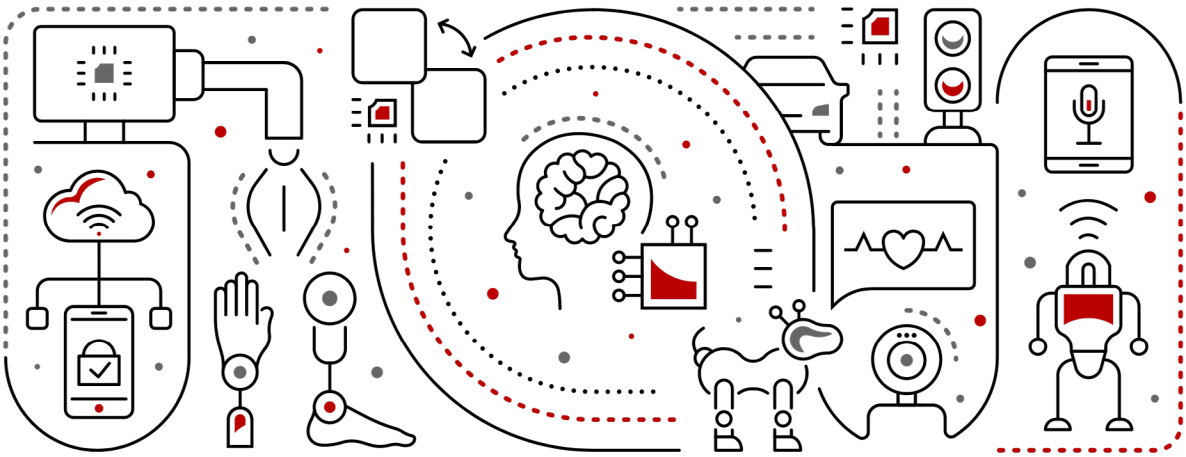

Artificial intelligence's (AI) dual nature is leading to a split response. Governments are seeking to promote the rapid development and use of AI systems, and businesses are enthusiastically embracing this project. At the same time, the media and the broader public are calling for greater controls on AI to ensure that it is developed and used in a safe and socially responsible manner, and governments are beginning to enact such requirements. Earlier this year, the Biden Administration issued an Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (E.O. 14110). Even more recently, the European Union decided to move forward with the first law (EU AI Act) to regulate AI comprehensively.

Businesses that develop and use AI are not insensitive to these trends. Eager to maintain public trust and to prepare for coming regulation, some companies are taking steps to make their AI systems safer, less biased, more protective of privacy, and otherwise more socially responsible. They refer to these new management activities as "responsible AI management," or "RAIM."

It is important that we not get too enamored of these business initiatives. Responsible AI management will neither make AI safe and equitable overnight, nor render government regulation and oversight redundant. But it is equally important that we understand RAIM, what it consists of, and what impact it has. Policymakers seeking to require businesses to use AI more responsibly need this information. Businesses seeking to achieve more socially responsible AI practices would benefit from it as well.

To date, the scholarship on responsible AI management has not kept pace with this growing need for information. There have been only a few empirical studies of RAIM,[1] and even fewer evaluations of RAIM practices. This report helps to fill this gap. It is part of a broader research effort, led by The Ohio State University Program on Data and Governance (PDG), to study private and public sector RAIM practices empirically. In 2018-2019, an Ohio State-led interdisciplinary group of scholars interviewed and surveyed leading businesses on their responsible AI management practices (then referred to as "data ethics" management). They published their findings in a report[2] and, most recently, an open-access book.[3]

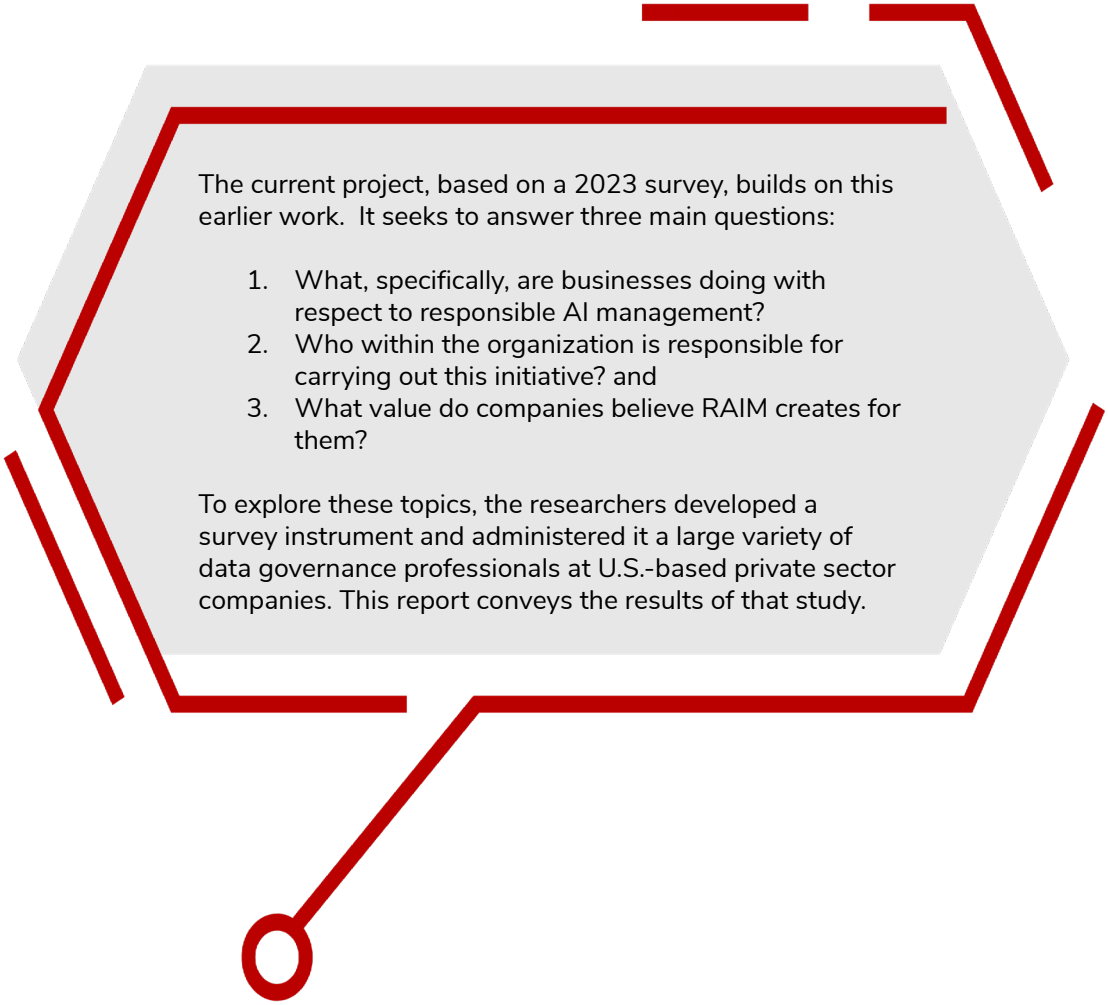

The current project, based on a 2023 survey, builds on this earlier work. It seeks to answer three main questions:

1. What, specifically, are businesses doing with respect to responsible AI management?
2. Who within the organization is responsible for carrying out this initiative? and
3. What value do companies believe RAIM creates for them?

To explore these topics, the researchers developed a survey instrument and administered it a large variety of data governance professionals at U.S.-based private sector companies. This report conveys the results of that study.

1 Centre for Information Policy Leadership, Building Accountable AI Programs: Mapping Emerging Best Practices to the CIPL Accountability Framework (2024); International Association of Privacy Professionals & FTI Consulting, Privacy and AI Governance Report: Privacy, Quo Vadis - Will You Lead the Way (2023); Cisco, Privacy's Growing Importance and Impact: Cisco 2023 Data Privacy Benchmark Study (2023); A. Renieris, D. Kiron, & S. Mills, To Be A Responsible Leader, Focus on Being Responsible (Sloan Management Review 2022); and E. Moss & J. Metcalf, Ethics Owners: A New Model of Organizational Responsibility in Data-Driven Technology Companies (Data & Society, 2020).
2 Hirsch, D. D., Bartley, T., Chandrasekaran, A., Norris, D., Parthasarathy, S., & Turner, P. N. Business data ethics: emerging trends in the governance of advanced analytics and AI (2021). https://moritzlaw.osu.edu/sites/default/files/2023-05/Final%20Report_10.15.21.pdf.
3 Hirsch, D. D., Bartley, T., Chandrasekaran, A., Norris, D., Parthasarathy, S., & Turner, P. N. Business Data Ethics: Emerging Models for Governing AI and Advanced Analytics (2024). https://link.springer.com/book/10.1007/978-3-031-21491-2.

METHODOLOGY

Drawing on prior work on business data ethics management (Hirsch et al., 2021), the researchers from the Ohio State University Program on Data and Governance designed and administered a survey that explored what U.S. businesses are doing with respect to responsible AI management (RAIM), and what value these businesses receive from such activities.

The International Association of Privacy Professionals (IAPP) provided early feedback on the survey design, assisted greatly with distributing the survey instrument, and was a particularly valuable partner throughout the implementation of this study. Cisco funded the research through an unrestricted gift.

It is often hard for academic researchers to identify corporate managers and get them to participate in a study. This task is particularly difficult for researchers studying emerging areas, such as responsible AI management, for which no clear corporate role exists as of yet.

01

Survey

To address this challenge, we distributed the survey to a wide variety of data governance professionals who were likely to know about, and be able to assess the value of, their company's responsible AI management activities. Professional organizations for privacy and other data governance professionals served as the primary means of reaching the target audience.

02

Distribution

As a collaborator on the project, the IAPP emailed the survey to 35,000 U.S.-based recipients of its Daily Dashboard. The Future of Privacy Forum (200+ members), Centre for Information Policy Leadership (85+ members), and Information Accountability Foundation (50+ members), also distributed the survey to their members. Subsequently, the researchers shared the survey via posts to LinkedIn groups connected to AI management, data governance, and AI ethics/responsible AI,[5] and via email to more than a hundred individuals whose LinkedIn profiles indicated they worked in RAIM or were high-level data governance professionals likely to have knowledge of their company's RAIM practices.

4 The names of these groups were: "AI Ethics and Values, Responsible AI," "AI for leaders," "Artificial Intelligence, Deep Learning, Machine Learning," and "Governance, Risk and Compliance Management (GRC)."

03

Screening

We provided these individuals with a clear definition of AI[5] and screened out those respondents who said that their company did not use AI as defined. We then provided a description of responsible AI management[6] and asked the survey respondents to self-select into the survey if they:

1. Had knowledge of their company's responsible AI management activities, and;

2. Had the ability to place a value on those activities.

To control for differences between regulatory and policy environments, we limited respondents to those whose company was primarily based in the United States. All survey responses were fully de-identified. The Ohio State Institutional Review Board approved the research design through an expedited review.

04

Response Rate

Participant recruitment took place from February through July of 2023. In all, respondents provided 75 completed or nearly completed surveys that met our selection criteria. This very low response rate may indicate that many busy data governance professionals did not have time to fill out a 20-25 minute survey. It could also mean that few companies today have meaningful RAIM programs in place, and so few individuals were in a position to describe and value them.[7] If sample members with less developed or non-existent RAIM programs were less inclined to answer the survey, the low response rate could further suggest selection bias in favor of those companies with the most developed RAIM programs. This possible selection bias, together with the fact that larger and more sophisticated companies were disproportionately represented in the survey sample, could mean that the survey results present a more optimistic view of responsible AI management than a more representative survey would reveal.

This report discusses the main findings of the study, providing a highlight of the descriptive statistics (frequencies, percentages, mean, etc.) and of other exploratory analyses. Given its nature, the data was analyzed using a multiplicity of non-parametric analyses to identify differences in the distribution of responses, as well as correlations with other variables.[8] In each case the pertinent analytical technique and results are indicated in the footnotes section.

5 The survey defined "AI" as "The advanced computational analysis of large sources of data or information, possibly by autonomous or semi-autonomous means, used primarily to forecast future outcomes, make predictions, guide decision-making, or generate recommendations. AI goes beyond simple descriptive analyses of data which are commonly used to gain business insights, and includes techniques such as data/text mining, machine learning, predictive analytics, natural language processing, neural networks, and other such complex quantitative methods."

6 The survey began with a consent to participate in research. The first paragraph following the consent stated that "This survey seeks information on responsible artificial intelligence (AI) management efforts at your company. The use of AI in its many forms can bring numerous benefits to business and contribute to the broader social good, but it can also raise a number of ethical dilemmas and risks. The survey asks about how you and your company manage these risks, and the value your company places on responsible AI management activities."

7 We did not collect the data we would need to assess the reasons that so many survey recipients did not participate, so our thoughts on this point are notional rather than grounded in data.

8 Tests used included Mann Whitney U-test, Wilcoxon signed-rank test, point biserial test, rank biserial test, and Spearman's correlation.

RESPONDENTS

Respondents mainly came from companies with 1,000 or more employees (79.7%) and with a revenue of $10,000,000 or more in the previous year (87.1%). Respondents also represented a variety of business sectors[9] including information technology, communication services, financial, consumer staples, consumer discretionary, industrial, healthcare, utilities, and real estate (Figure 1). The large size of the responding companies may be a function of the channels used to distribute the survey. But it also may suggest that larger companies are doing more with respect to RAIM than small- or mid-sized ones and so were more likely to respond to the survey. The wide range of sectors represented suggests that the use of AI, and the resulting need for AI governance, is present in many industries not traditionally identified as being in the "tech" sector. That said, the information technology sector represented 44.6% of the survey respondents, potentially reflecting this sector's greater focus on responsible AI management. When compared to the distribution of sectors in the IAPP-EY Privacy Governance Report,[10] which employed a similar sample frame to that used in our study, we see the respondents coming from the information technology sector greatly surpass the 13% of respondents from the tech/telecommunication's sector in the IAPP-EY report.

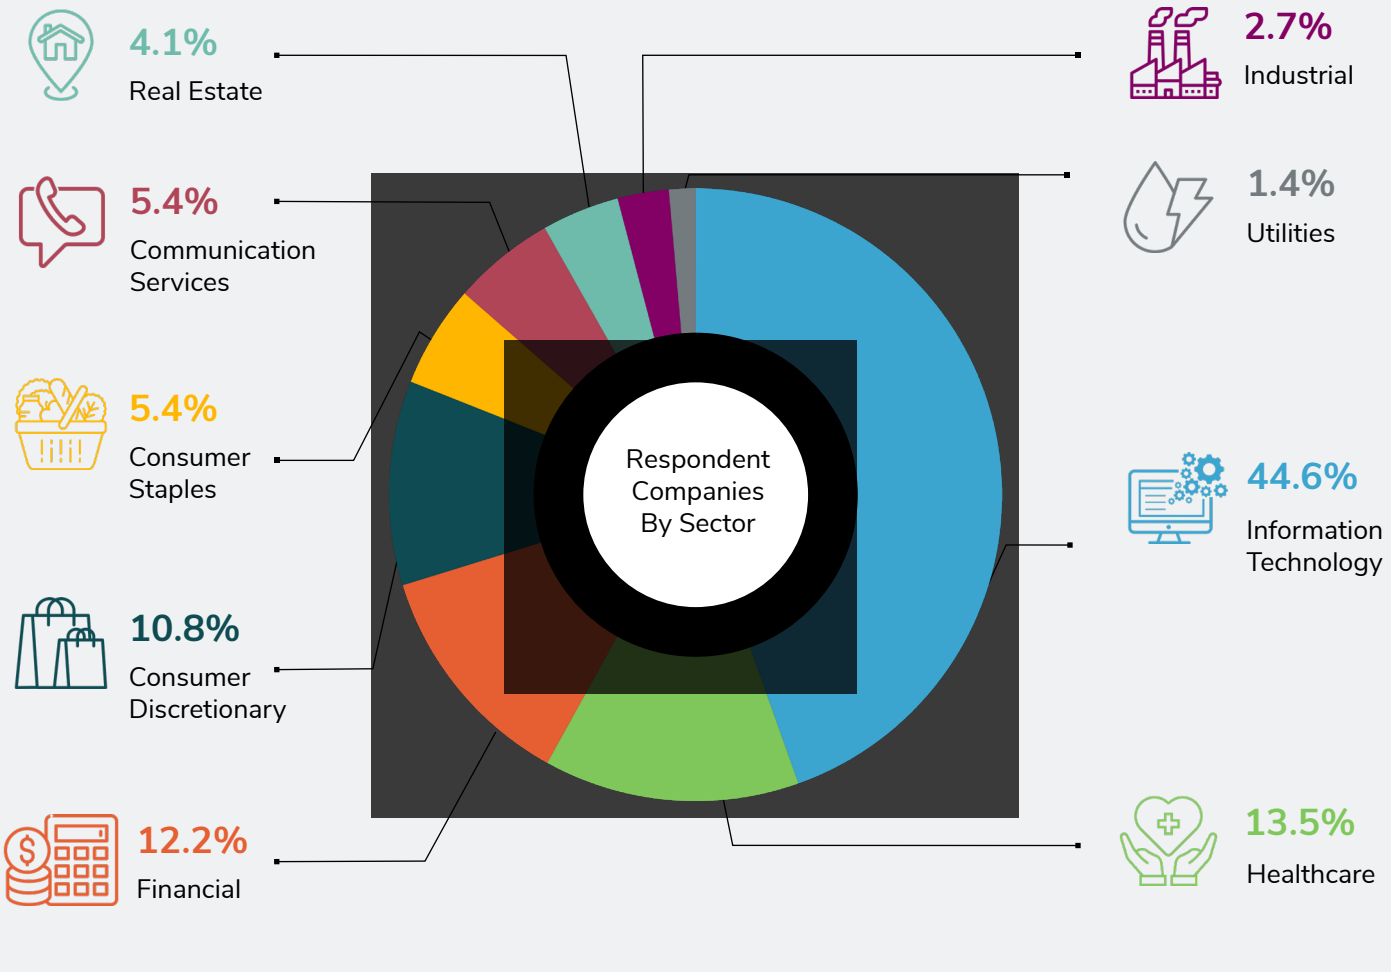

Figure 1: *Respondent companies by sector.*

9 Sectors were taken from *The Global Industry Classification Standard (GICS®)* – https://www.msci.com/our-solutions/indexes/gics
10 IAPP-EY Privacy Governance Report 2023, available at https://iapp.org/resources/article/privacy-governance-report/

RAIM ATTITUDES & IMPLEMENTATION

68%
of survey respondents said RAIM was either **important or extremely important** to their company (Figure 3).

Nearly 90%
either **"Agreed" or "Strongly Agreed"** with the idea that companies should make a **"meaningful investment"** in RAIM, (Figure 2).

A large majority of the survey respondents (68%) said that RAIM was either important or extremely important to their company. (Figure 3). Nearly 90% either "Agreed" or "Strongly Agreed" with the idea that companies should make a "meaningful investment" in RAIM. (Figure 2). The favorable view toward RAIM investment did not vary substantially between respondents who reported that RAIM provided high overall value to their company, and those who reported lower overall value. These data suggest an overwhelming, generalized belief – among respondents at least – in the value of RAIM and in the need to invest in it.

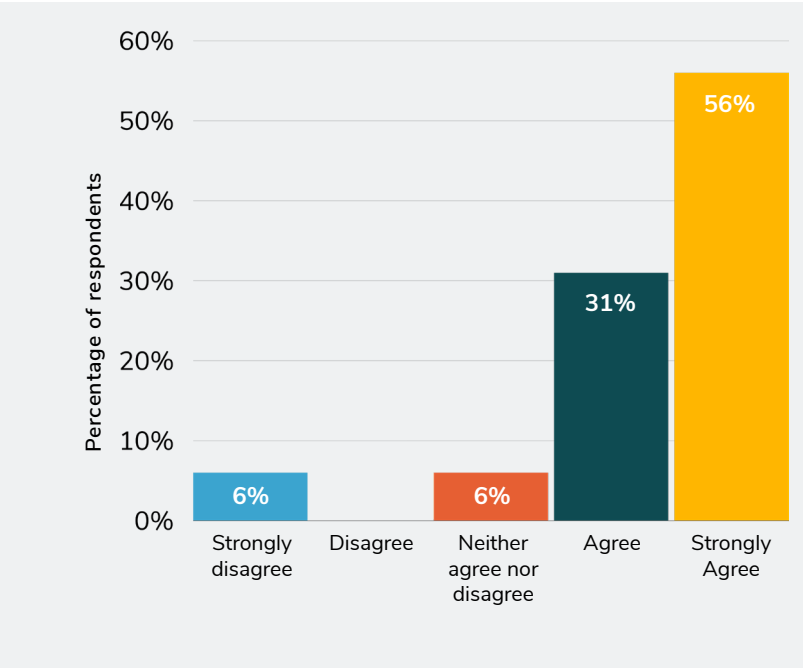

Figure 2:
Extent to which respondents believe that companies should make meaningful investment in RAIM.

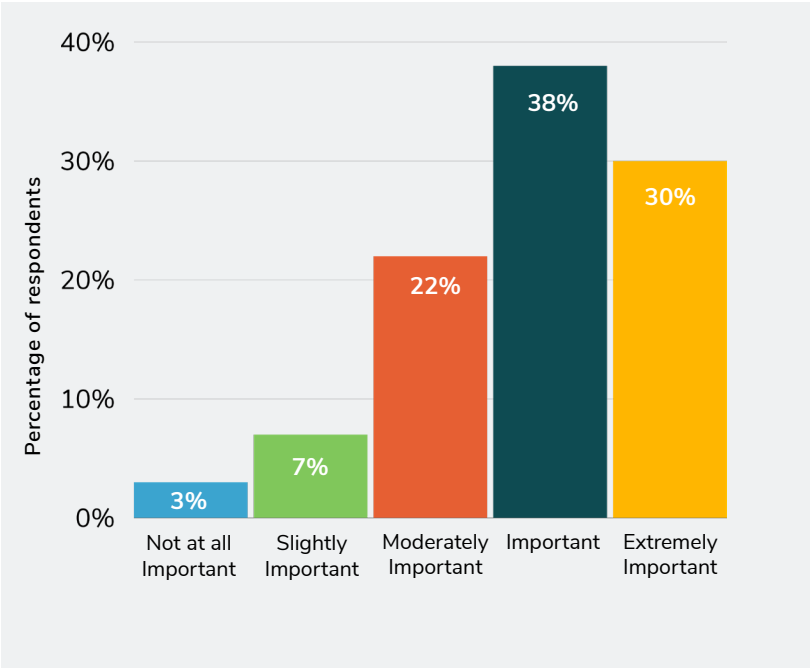

Figure 3:
Importance of responsible AI management to respondent companies.

Implementation lagged enthusiasm, however. Only 29% of respondents stated that their companies had implemented RAIM to a "large extent" or to a "very large extent," while 71% indicated that they had implemented them "somewhat," "very little," or "not at all." The most common response (44%) was that the company had implemented RAIM "somewhat" (Figure 4).

Figure 4:
Extent to which RAIM practices have been implemented at your company.

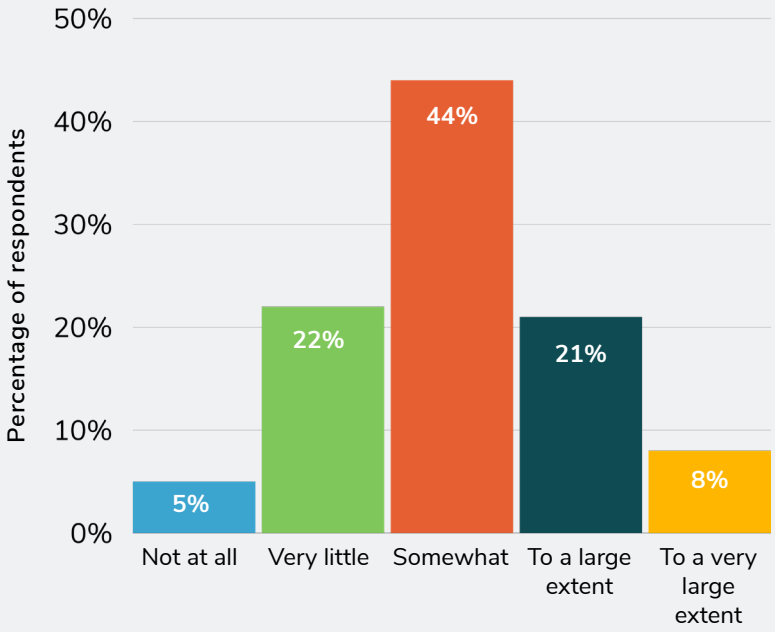

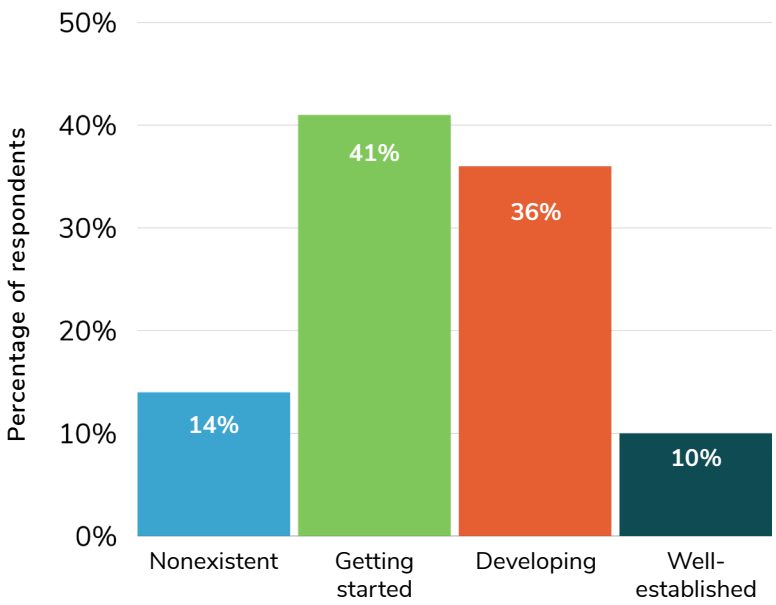

The survey also asked about the maturity level of the RAIM programs at the responding companies. More than 75% of respondents described their company's RAIM program as either "New/Young/Getting Started" or "Developing/Growing," while less than 10% described it as "Mature/Well-established", and nearly 14% of respondents selected "Non-existent/Not Yet Implemented" (Figure 5).

Figure 5:
Companies' RAIM program maturity level

Nearly two-thirds of respondents described their company's process for making responsible AI judgment calls as "ad hoc" rather than "systematic" (Figure 6). This is particularly striking given that the responding companies were quite large and, if indeed selection bias shaped the survey results, would be those likely to have the most advanced RAIM programs. Even large, well-resourced companies tended to use "ad hoc" processes for making AI ethics judgment calls.[11] This further indicates the immaturity of RAIM today.

Figure 6:
Is your company's process for making AI ethics judgment calls ad hoc or systematic?

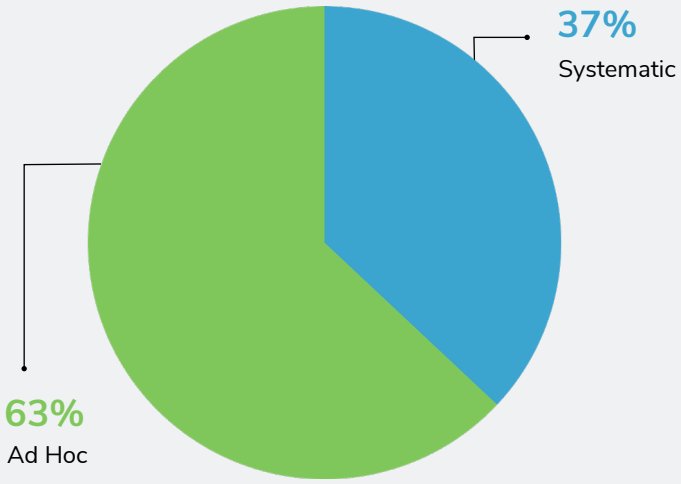

The gap between companies' expressed belief in the importance of RAIM, and their middling implementation of it, could reflect the lack of trained professionals with the skill sets required to plan and institute such programs. Alternatively, it could be a function of the time that it takes for organizations to translate intentions into action. Over 70% of respondents anticipated that their companies' investment in RAIM would grow either "somewhat" or "by a lot" over the next year (Figure 7).

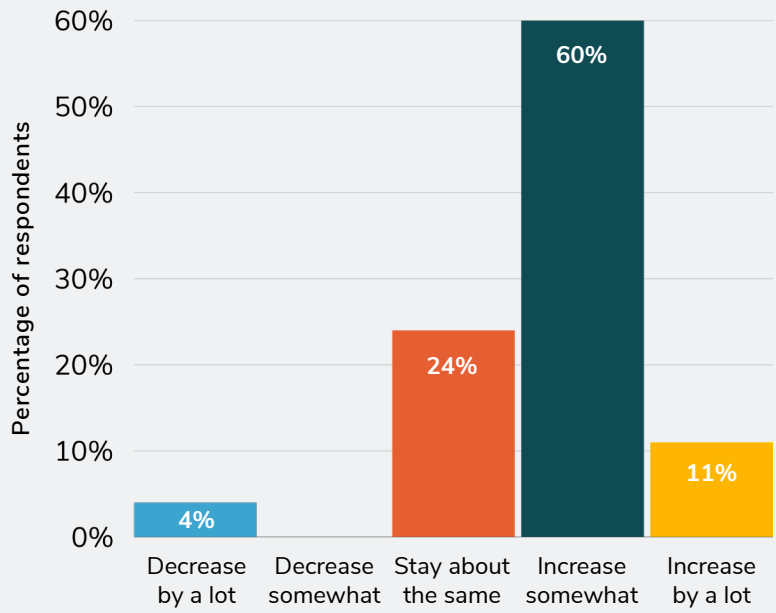

Most of the remaining respondents stated that their RAIM investment would stay the same, and only a small number indicated that it would decrease. Companies that had done more with respect to RAIM implementation expected a higher increase in future investments than those who had done less. This could indicate that, as an organization starts to implement RAIM, it gains a more realistic view of what is needed and a greater appetite for increased investment.[12]

Figure 7:
Respondents' expectation for their company's investment in RAIM over the next year.

11 Larger companies, measured by revenue and number of employees, were as likely to indicate they use ad hoc processes than mid-sized and smaller companies (U test for revenue p=0.13, U test for number of employees p=0.13)
12 Companies that have a developing or well-established plan tend to expect a higher increase in future investments than those that are getting started or have a nonexistent program. (rs=0.25, p=0.09). In the same way, those that have implemented systematic processes tend to expect a higher increase in future investments (p=0.03)

KEY FEATURES OF A RAIM PROGRAM

When organizations say they are pursuing responsible AI management what, specifically, do they mean? To answer this question, we began by reviewing recent literature on data ethics and responsible AI management to identify the important features of RAIM programs.

We selected 14 commonly-reported features and asked respondents whether their company engaged in each. We further asked who in the organization was responsible for the company's RAIM activities. This section conveys the survey results on these points. It details what companies are doing with respect to responsible AI management, and who is involved in the effort. A subsequent section reports the value that the respondents said these activities create for the company and teases out some of the relationships between actions taken and value produced. Figure 8 identifies the main components of respondent companies' RAIM programs, and the percentage of companies that adopt them.

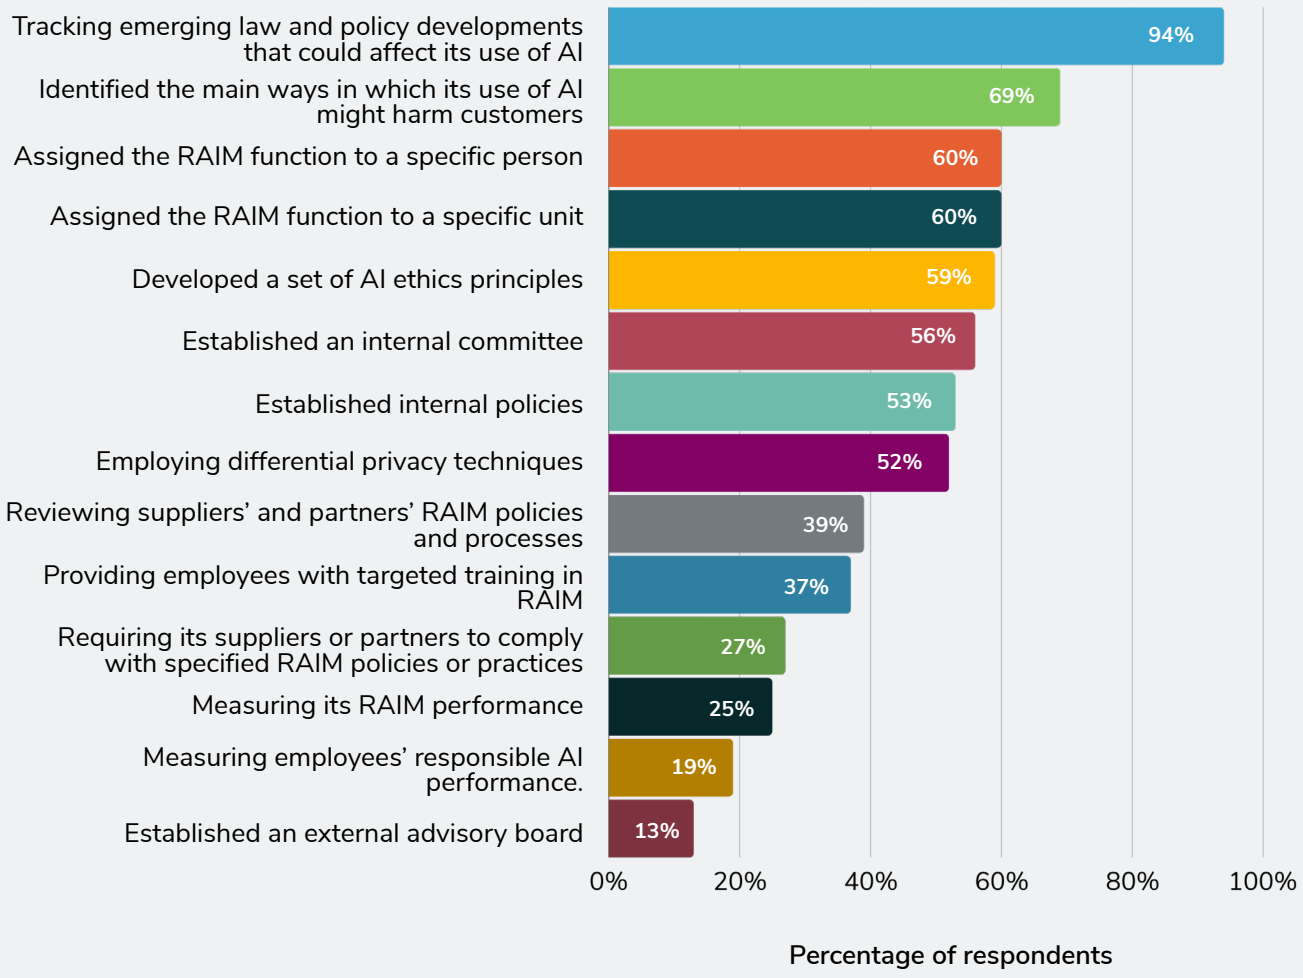

Figure 8:
Responsible AI management activities companies are engaging In.

As Figure 8 shows, 94% of respondents said that their organization tracked law and policy developments that could affect its use of AI. Sixty-nine percent said that their organization identified the ways that the company's use of AI might harm customers or stakeholders. Sixty percent indicated that their organizations had assigned the RAIM function to a specific person, and a similar percentage said that they have assigned it to a specific unit. Fifty-six percent of companies established an internal committee (sometimes called an "AI ethics committee") to make decisions about what it means for the company to use AI responsibly. In contrast to the popularity of internal AI ethics committees, only 13% of respondents reported the creation of an external advisory board to provide input on AI ethics issues.

Fifty-nine percent of respondents said that their organization had defined a set of ethical principles to guide its use of AI, and 53% said that they had established internal policies on responsible AI for frontline data scientists. However, only 39% of respondents said that their company reviewed suppliers' and partners' responsible AI practices and policies, and only 27% required suppliers and partners to comply with the company's own responsible AI policies or practices. In contrast to the 53% who said that they had established internal policies on responsible AI, only 37% of responding companies said that their organization provided their employees with targeted training in this area. Fifty-two percent reported employing differential privacy to protect the identities of individuals in the training data.

Finally, most responding companies do not measure either their own, or their employees', responsible AI performance. Only 25% of companies measure their own responsible AI performance, and only 19% look at how well an employee meets the company's responsible AI goals when evaluating that person's performance and/or setting their compensation.

These data suggest several trends.

01 Tracking law and policy developments (94%) is a way of assessing regulatory risk to the organization. Identifying the ways in which the company's use of AI could harm others (69%) assesses risk to those outside the organization. Taken together, the popularity of these two actions suggests an interest in risk assessment.

02 Companies seem to be putting an emphasis on building a RAIM management structure. We see this in the number of companies that have appointed a responsible person (60%) and/or created an AI ethics committee (56%).

03 The number of companies adopting AI ethics principles (59%) and responsible AI policies (53%) suggest an interest in developing RAIM policies, although this is not yet matched by a similar commitment to providing their own employees with targeted RAIM training (37%) or to reviewing their suppliers' and partners' RAIM policies (39%).

Finally, we see only limited efforts towards evaluating their own (25%) or their employees (19%) RAIM performance. This is an important gap. Measuring corporate performance in a given area can be critical to improving that performance. Similarly, the criteria by which a company evaluates employee performance communicate what a company values and can be an important lever for motivating and changing employee behavior and priorities. The importance of such practices, and the fact that most of the companies in our sample are not yet taking advantage of them, suggests that companies could readily improve their RAIM program by measuring organizational and employee responsible AI performance and then taking steps to improve that performance. A barrier to doing so may be the lack of reliable metrics for conducting such evaluations. More research on how some companies currently conduct such evaluations, and the metrics that they are using, as well as more theoretical explorations of the metrics that such companies should be using, would be useful here.

WHO IS RESPONSIBLE FOR AI MANAGEMENT?

In addition to understanding what organizations were doing with respect to responsible AI management, we also wanted to get a grasp on who was doing it. As reported above, the majority of responding companies assigned responsibility for the RAIM function to one or more responsible persons or units. Whom did they select for this role?

Figure 9 shows great variability in the title of the person or people responsible for RAIM. The survey presented ten options ranging from data analytics manager, to privacy officer, to legal counsel, to C-suite executive, and beyond. Each of these options received some affirmative responses.[13] Even this did not even cover the map. Nine percent of respondents selected "other."

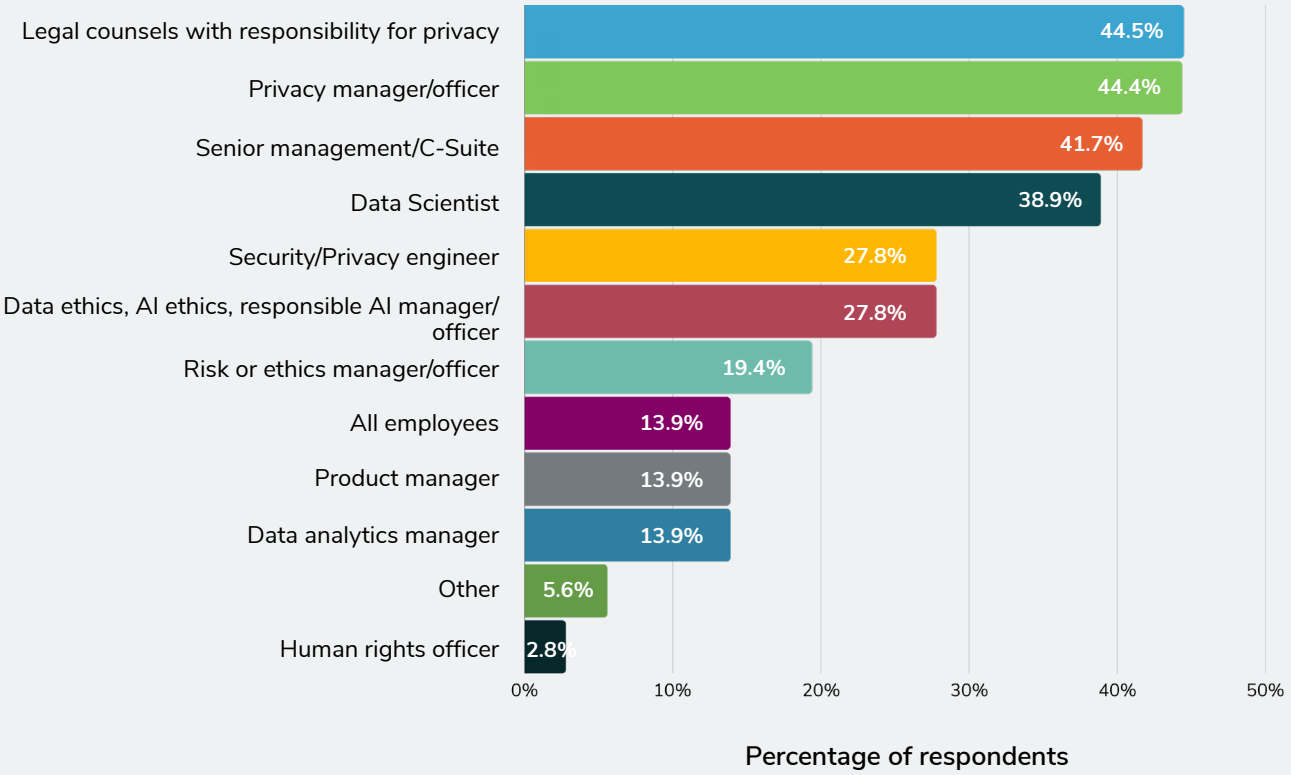

Figure 9:
Which of the following titles best matches the title(s) of the individual or individuals to whom your organization has assigned the RAIM function?

13 This was presented as a multiple choice question and respondents could select more than one position title. That is why the categories total to well more than 100%.

The percentages in Figure 9 total to well more than 100 percent. This is because the survey allowed respondents to identify more than one title for the "individual or individuals" to whom the organization had assigned the RAIM function, and some respondents identified two or more positions as holding this responsibility. As Figure 9 shows, "Privacy Manager" and "Legal Counsel with Responsibility for Privacy" led the pack with just over 44% of respondents indicating that they had assigned at least some RAIM responsibility to this position. "Senior Manager" and "Legal Counsel With Responsibility for Privacy" followed closely behind at 41.7%, followed by "Data Scientist" at 38%. Roughly twenty-eight percent allocated at least some of the responsibility to an AI Ethics or Responsible AI officer, an indication that more than a quarter of responding companies have created such a position.

Respondents also shared the units where the individuals with assigned RAIM responsibility were located. Responses suggest that there is not a single way to set the management structure for the RAIM function, as it is possible to find individuals playing a similar role, but embedded in different units or departments, across the participating companies. For example, the "Data ethics, AI ethics or responsible AI manager" can be found in a dedicated ethical or responsible AI unit in some cases, while in others this position can be part of business units, privacy, legal IT, or other units.

It is further enlightening to group the positions by similarity of expertise. For example, we can combine "Privacy Manager/Officer", "Security/Privacy Engineer" and "Legal Counsel with Responsibility for Privacy" into a privacy-focused category, and "Data Scientist" and "Data Analytics Manager" into an analytics-oriented one. Grouping the positions in this way (Figure 10) shows that companies are most likely to assign RAIM responsibility to those with expertise in privacy. The organizations that distributed the survey (IAPP, FPF, IAF, CIPL) primarily cater to privacy professionals. Selection bias may accordingly have influenced the results in Figure 10. Still, it stands to reason that companies would task those with expertise in data governance with the responsibility for AI governance.

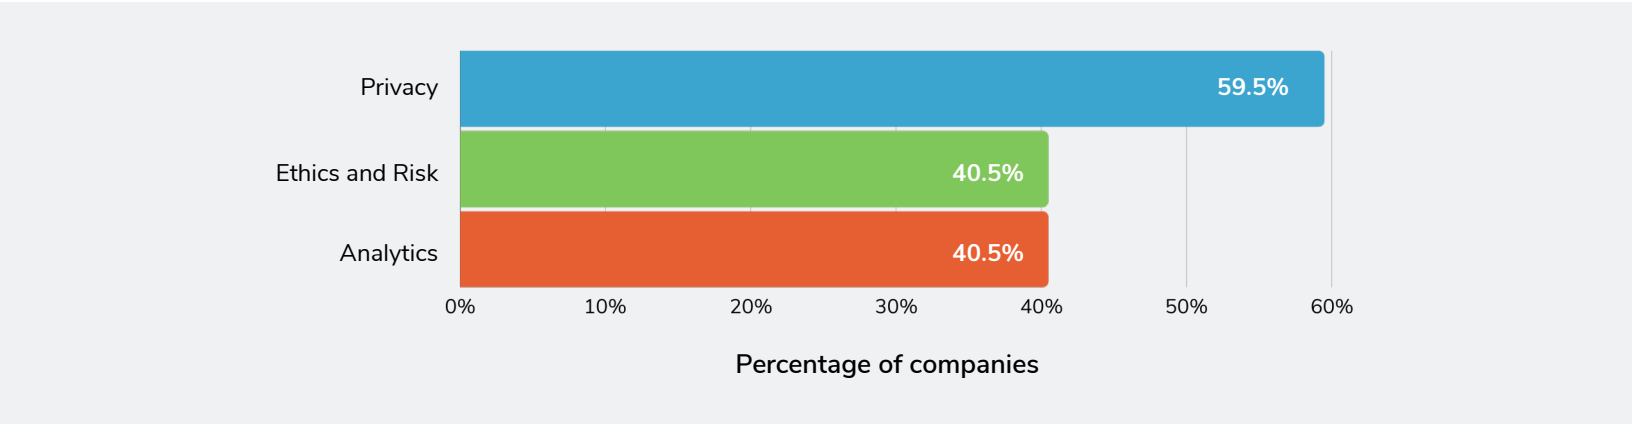

Figure 10:
Percentage of companies that assigned responsibility to a person in area of expertise.

The overall impression remains one of considerable diversity in the personnel involved in responsible AI management. This broad mix of professionals could mean that companies require various types of expertise – legal, data science, privacy, governance, etc. – to govern AI effectively. Additional research will be required to confirm this. The strong representation of Senior Officials and C-suite executives could indicate that RAIM has strategic importance, can have substantial impacts on the company, and so warrants the attention of a senior executive.

RAIM PROGRAMS CREATE VALUE

Having asked about RAIM activities and the people responsible for them, the survey turned to the value that RAIM creates for companies that pursue it.

The results were striking. Responsible AI management programs create substantial value for the businesses that implement them (Figure 11). Companies with more developed RAIM programs report gaining more value from this management activity, on average (Figure 12).

Figures 11 and 12 are based on respondents' subjective estimates of their RAIM program's value, not an assessment of these programs' actual value or a comparison of this value to the cost of standing up a RAIM program. Still, the results are impressive. As Figure 11 shows, all respondents reported that their company gets at least some value from their RAIM programs; none said that they get no value. Close to 40% reported that their company gets "a lot" or a "great deal" of value from their RAIM programs. In the eyes of those well-positioned to assess it, RAIM programs do create value and can, when more fully developed, generate more of it on average. Responsible AI management can be a source of business value.

Nearly 40%
reported that their company gets **"a lot" or a "great deal" of value from their RAIM programs.**

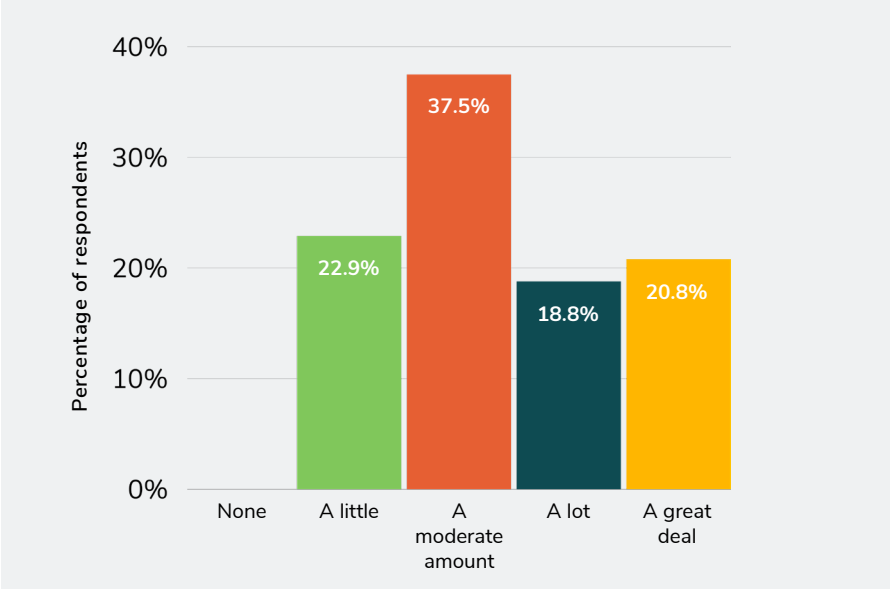

Figure 11: *Overall value from investment in RAIM.*

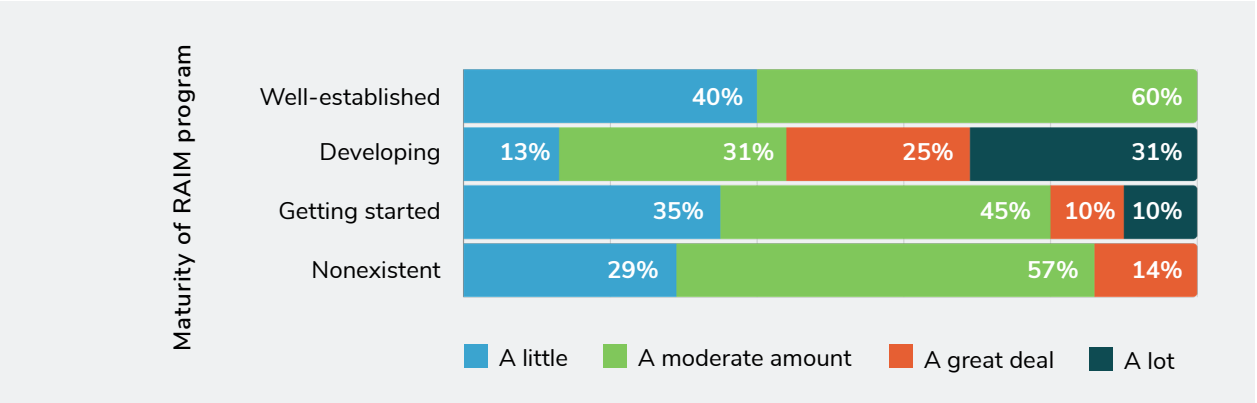

Figure 12: *Reported value for companies at each RAIM program maturity level.*

The survey further explored the types of value that RAIM programs create. To examine this, we reviewed the literature and identified six core ways in which RAIM programs might provide value: (1) Improving the company's reputation for trustworthiness; (2) Reduction of regulatory and legal risk; (3) Improving product quality; (4) Improving employee relations; (5) Providing competitive advantage; and (6) Supporting corporate values. We then asked respondents to report the level of value that their company's RAIM activities created in each of these areas. Figure 13 presents their responses.

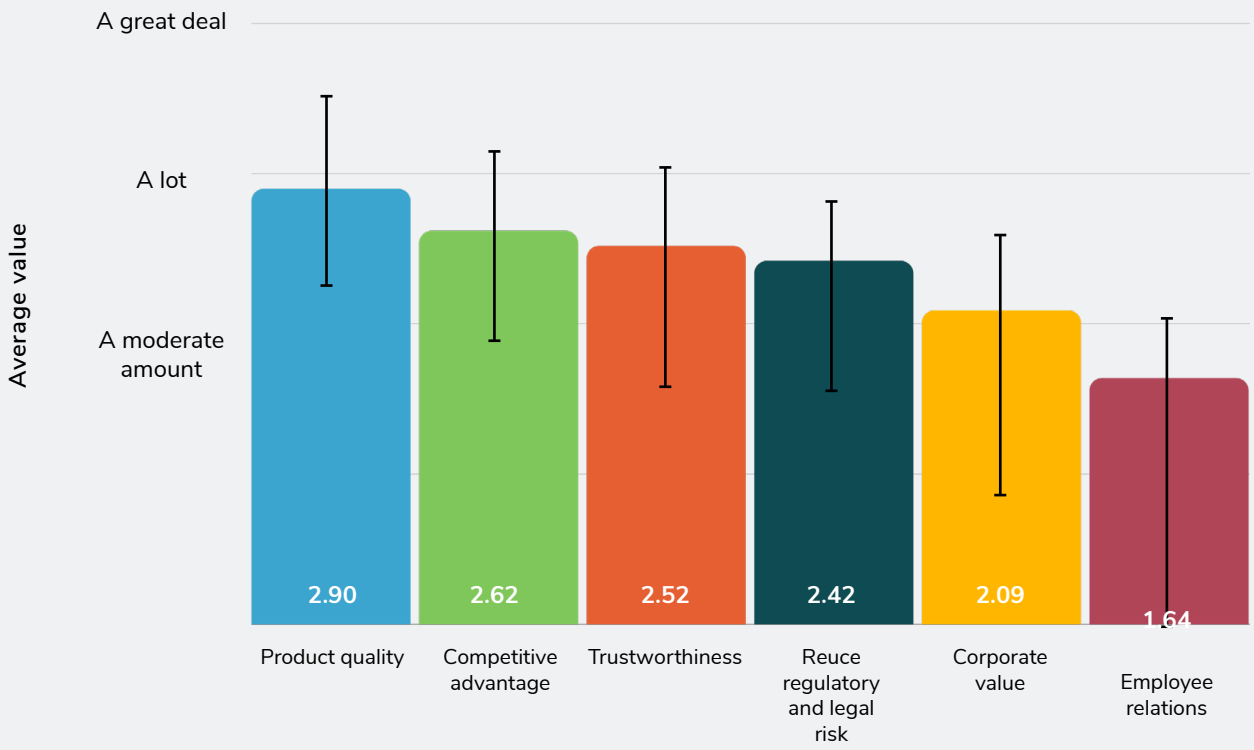

Figure 13:
Average level of value reported for each value category.[15]

The data suggest that, while RAIM programs do allow companies to define and communicate their corporate values better, they do far more than this. They improve product quality, make the business more trustworthy, reduce regulatory and legal risk, and enhance employee relations. Each of these categories is a source of strategic value for a company. RAIM programs serve a strategic function, in addition to allowing a company better to attain its corporate values. Indeed, according to the respondents, RAIM's greatest source of value lies in improvements to product quality, with increases in competitive advantage, trustworthiness and risk reduction close behind. RAIM provides more value in each of these areas than it does by enabling a company better to define and achieve its corporate values.

14 Relationship between maturity of the RAIM program, and the overall level of value is positive and significant (rs=.50, p= <0.01)
15 Black whiskers represent standard deviation.

Figure 14 breaks each of the broader value categories into its component parts, and so allows us to look even a bit deeper into the value that RAIM creates for the companies that practice it. To begin with, Figure 14 indicates that RAIM provides competitive advantage more by increasing innovation and improving revenues, than by decreasing costs.[16] Similarly, RAIM improves product quality more by improving AI products' accuracy and alignment with consumer expectations, than by reducing defects. To date, most have viewed RAIM as a way of reducing risk and liability. These data suggest that RAIM also pushes companies to innovate and enhance their AI products and processes ways that increase revenue and that this is, in fact, the larger source of value and of competitive advantage.[17]

These data are hardly definitive. The differences in the respondents' ratings of the different value components are slight[18] and do not allow for hard conclusions. However, they suggest an important, new way of thinking about responsible AI management – as a source of value and competitiveness, and not just a way of mitigating risks and costs. This is an important area for further research.

Figure 14 further suggests that, when it comes to "reducing regulatory and legal risk," RAIM programs provide the most value by making companies "better prepared for regulation" and the least by "reducing the chance of regulation."[19] This could mean that, given the rapid growth in AI law and policy, respondents see increased regulation of AI as inevitable and so focus more on preparing for these new laws than on pre-empting them.

16 In further support of this idea, "increase competitive advantage" was reported to provide consistently higher value than "increase revenue" and "decrease costs," while providing similar value to "increasing innovation." (Wilcoxon signed rank test with a significance above 95%). This suggests that respondents see competitive advantage less in terms of reduced costs or increased revenues and more in terms of enhanced innovation.

17 The respondents gave similar rankings to "increase competitive advantage," "increase innovation," and "AI products support consumer expectations," thereby suggesting that they saw these three sources of value as linked (Wilcoxon signed rank test with a significance above 95%).

18 Particularly the difference between "increasing innovation" and "decreasing costs."

19 Respondents ranked "reduce liability" and "prepare for regulation" consistently higher than "reduce the chance for regulation" and "influence regulation." (Wilcoxon signed rank test with a significance above 95%)

Figure 14:
Value categories broken down into component parts.

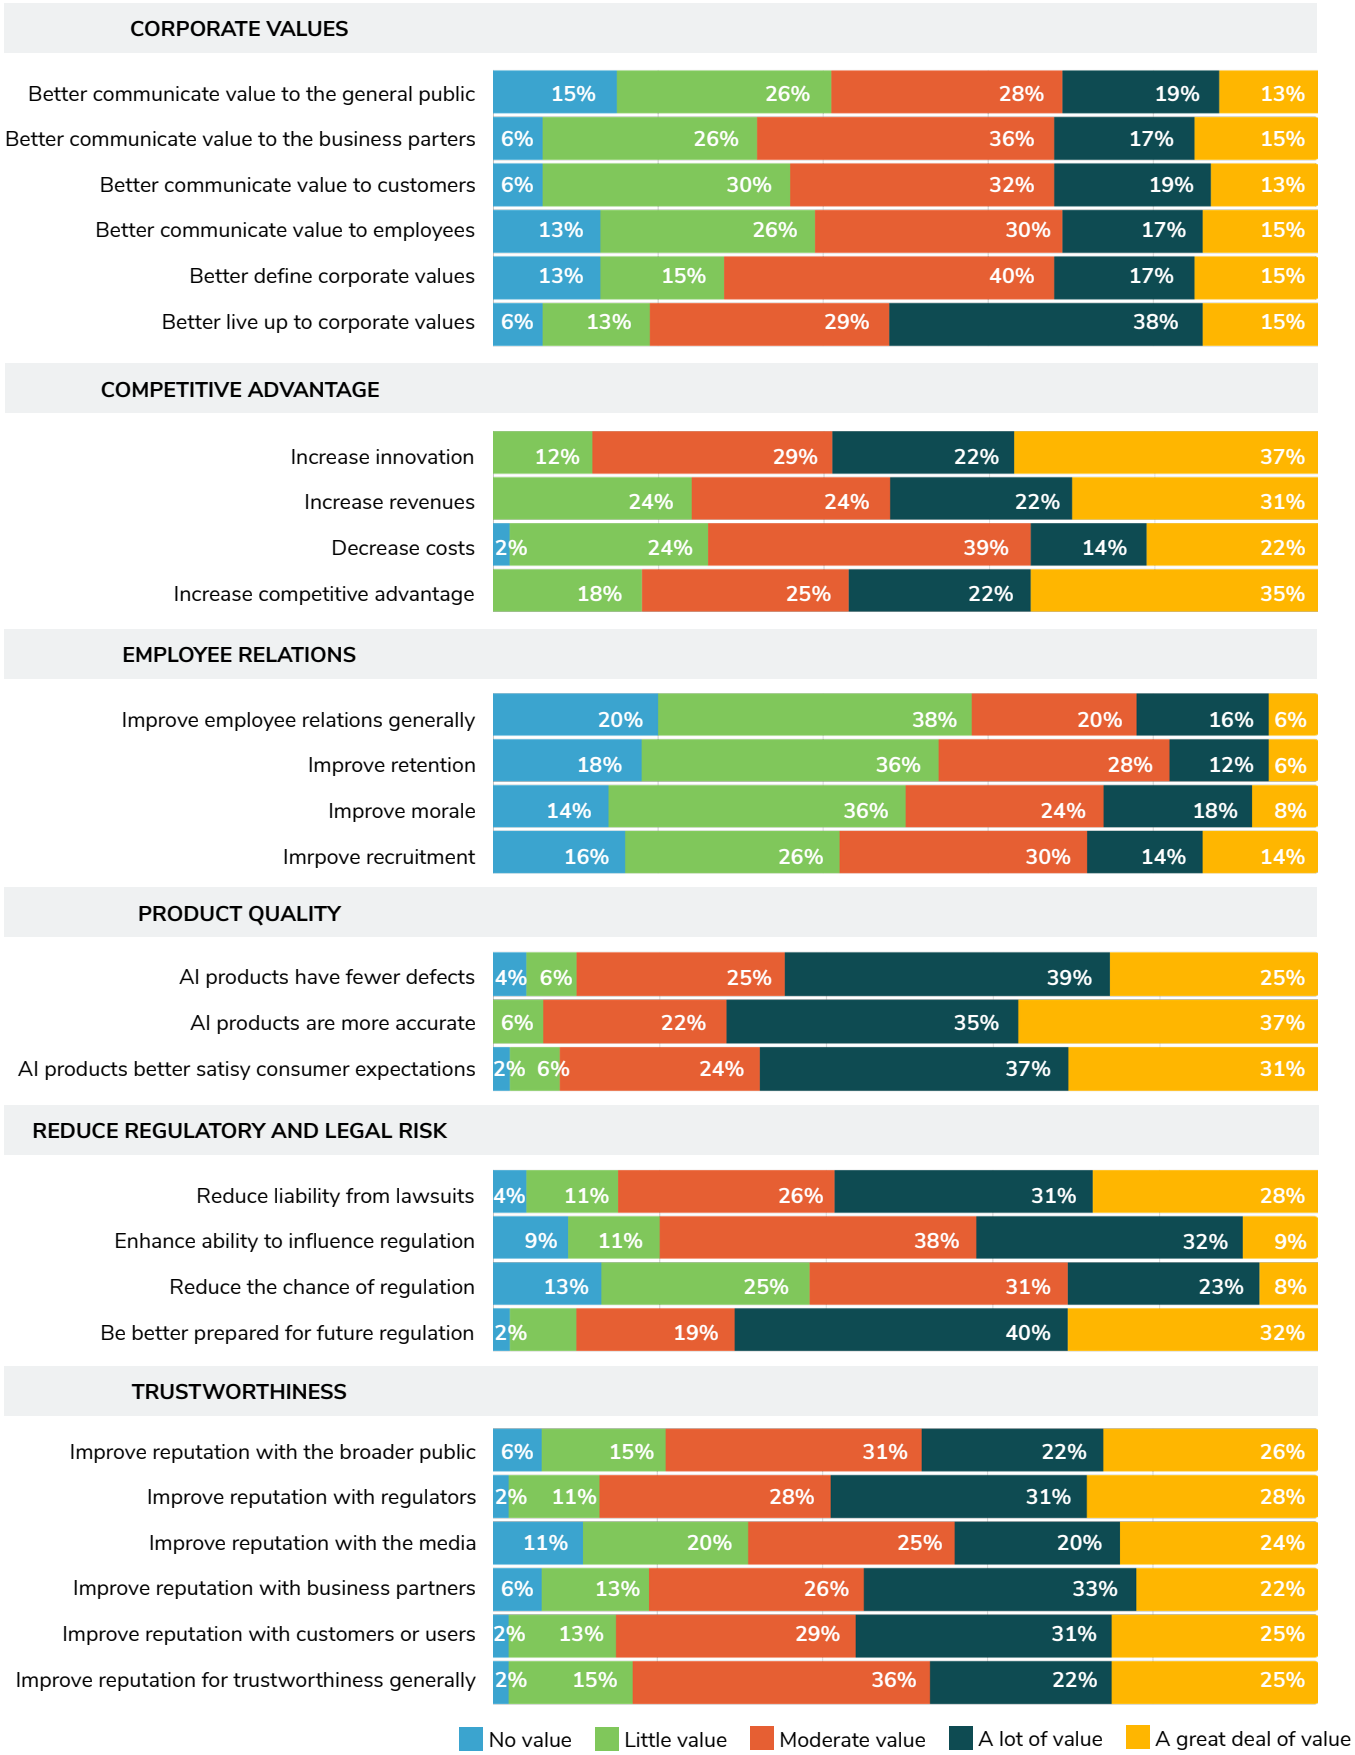

Finally, with respect to improving “trustworthiness,” respondents highlight RAIM’s impact on improving reputation among regulators and with consumers or users, and place less importance on its ability to improve reputation with the media, business partners or other stakeholders.

The findings on RAIM value are still exploratory. Further research is needed to develop a more refined understanding of the types of value that RAIM generates, and to move from measuring managers’ perceptions, to assessing RAIM programs’ actual outcomes and impacts.

RAIM ACTIVITIES DRIVING VALUE

The prior two sections discussed what RAIM consists of, and the value that it creates for organizations. This one explores the relationship between specific RAIM activities, and respondents' assessments of the value created in each of the seven general value categories.

These analyses begin to shed light on which RAIM activities generate business value, and which types of business value they create. In several instances, particular RAIM activities were significantly correlated with specific types of business value. Here we discuss only the most meaningful correlations.[20]

First, establishing internal policies and requiring suppliers to comply with specified policies seem to have a positive influence on trustworthiness. As Figure 15 shows, each of these activities is positively correlated to all the items included in trustworthiness. However, "establishing internal policies" seems to have a stronger impact on improving reputation with regulators, the media, and the broader public, while "requiring partners to comply with policy" has a stronger relationship to improving reputation with customers. Additional research will be required to confirm whether, and why, these relationships exist. For now, we hypothesize that companies that establish internal RAIM policies should be able to pursue RAIM more deliberately and consistently than those that utilize a more "seat of the pants" approach, and that this may produce greater stakeholder trust.

A company that not only establishes internal policies but also requires its suppliers to comply with them will be even more consistent and should be better protected against the reputational harm that a suppliers' irresponsible use of AI can inflict. This may be why such companies experience a greater increase in consumer trust than those that do not impose their RAIM policies on suppliers.

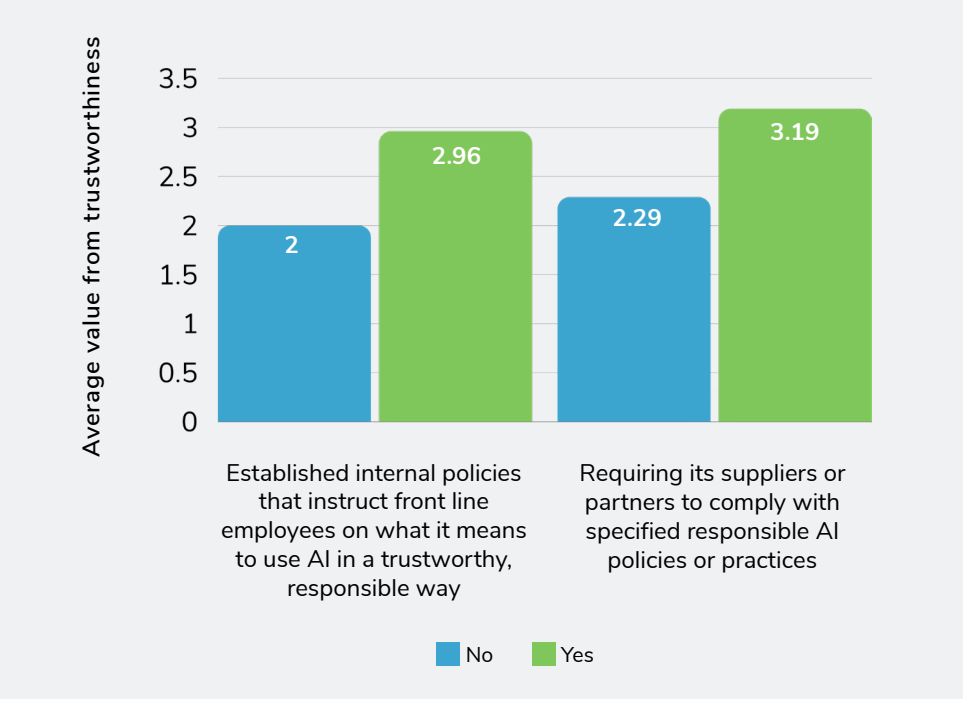

Figure 15:
Average perceived value from trustworthiness dimension by RAIM activities.

20 The relationships discussed here have at the same time a statistically significant t-test with a 90% confidence, and a present a mean difference equal or greater than .85, and significant point biserial correlation with a rho higher than .35.

Second, as Figure 16 indicates, the survey data show companies that identify potential harms that AI products could cause to customers and other stakeholders, as well as those that track the emerging law and policy, experience more improvements in product quality. This relationship seems to hold for every component of improving product quality. However, the strongest correlation is with achieving products with fewer defects. These data suggest that companies may increasingly view AI product or process "defects" not just in terms of inaccuracies or performance deficits, but also in terms of harms to others and legal violations. This would make responsible AI a new baseline for AI systems such that a departure from this standard counts as a defect.

The relationships illustrated in Figure 16 may further indicate that pro-actively identifying potential harms and tracking laws and policies puts a company in a position to prevent such harms and/or violations, and so to improve product quality. Again, further research will be required to establish these connections.

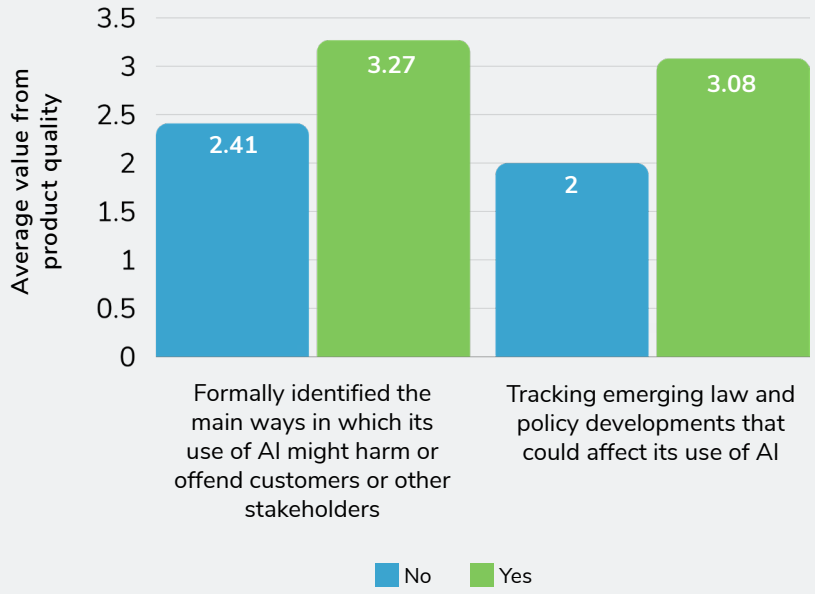

Figure 16:
Average perceived value from product quality dimension by RAIM activities.

Third, as shown in Figure 17, "providing employees with targeted training" and "measuring the company's RAIM performance" have the most positive impact on improving employee relations. This correlation seems to be driven in both cases by the value coming from improving morale and retention among existing employees, and not from improvements to recruitment of future ones or to employee relations more generally.

This distinction could result from the fact that employees trained in RAIM have more information about a company's commitment to this area than those whom the company is trying to recruit. If so, this would suggest potential value in sharing with recruits information about the company's RAIM activities and commitment

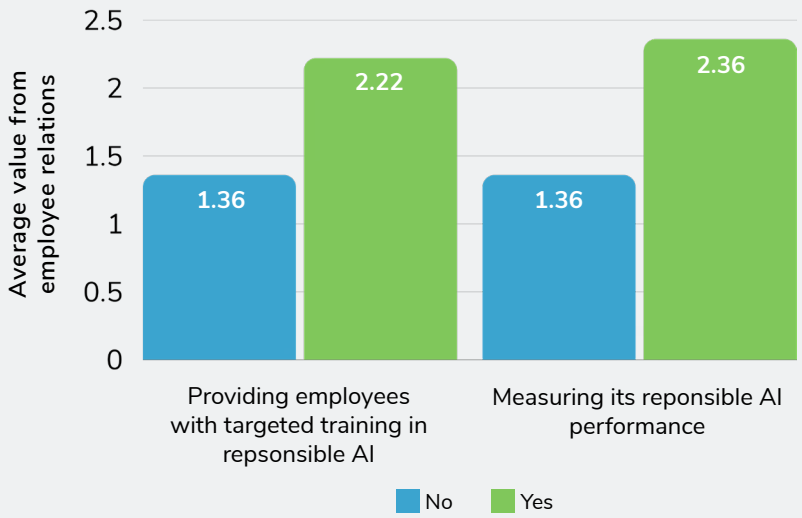

Figure 17:
Average perceived value from employee relations dimension by RAIM activities.

As shown in Figure 18, "requiring partners and suppliers to comply with RAIM policies" increases competitive advantage. This activity is correlated to every item in competitive advantage, although the strongest connection is to decreasing costs. This suggests that setting RAIM policies for the entire supply chain helps to reduce AI ethical or legal violations that can prove costly for a company even if committed by a supplier.

Finally, as shown in Figure 19, multiple activities seem to influence a company's ability better to live up to its corporate values. These are: formally identifying potential harms to customers and stakeholders, assigning the RAIM function to a specific person, training employees in RAIM, and measuring the company's RAIM performance. If we break "better achieve corporate values" down into its component parts, we can see even better how these various activities have this positive impact. Specifically, each of these activities is strongly correlated with communicating corporate values to employees, customers, partners and the general public. Measuring performance is correlated to the value of better living up to the corporate values. By contrast, none of the activities are significantly correlated with better defining corporate values. This could be because most companies have already defined their corporate values. They do not need RAIM for this. They need it for applying these values to their new AI operations, and so maintaining and fulfilling their values in the AI age.

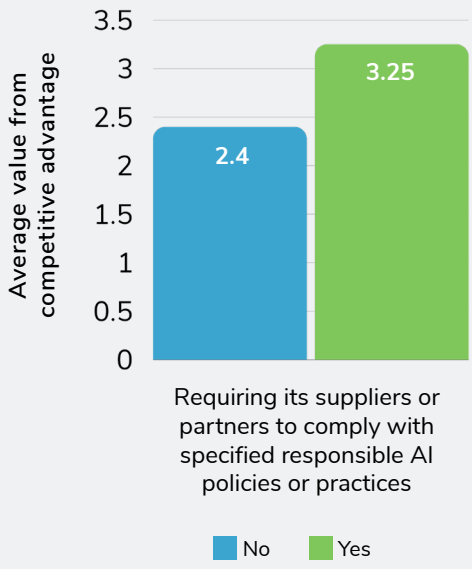

Figure 18:
Average perceived value from competitive advantage dimension by RAIM activities.

Figure 19:
Average perceived improvement in better living by corporate values by RAIM activities.

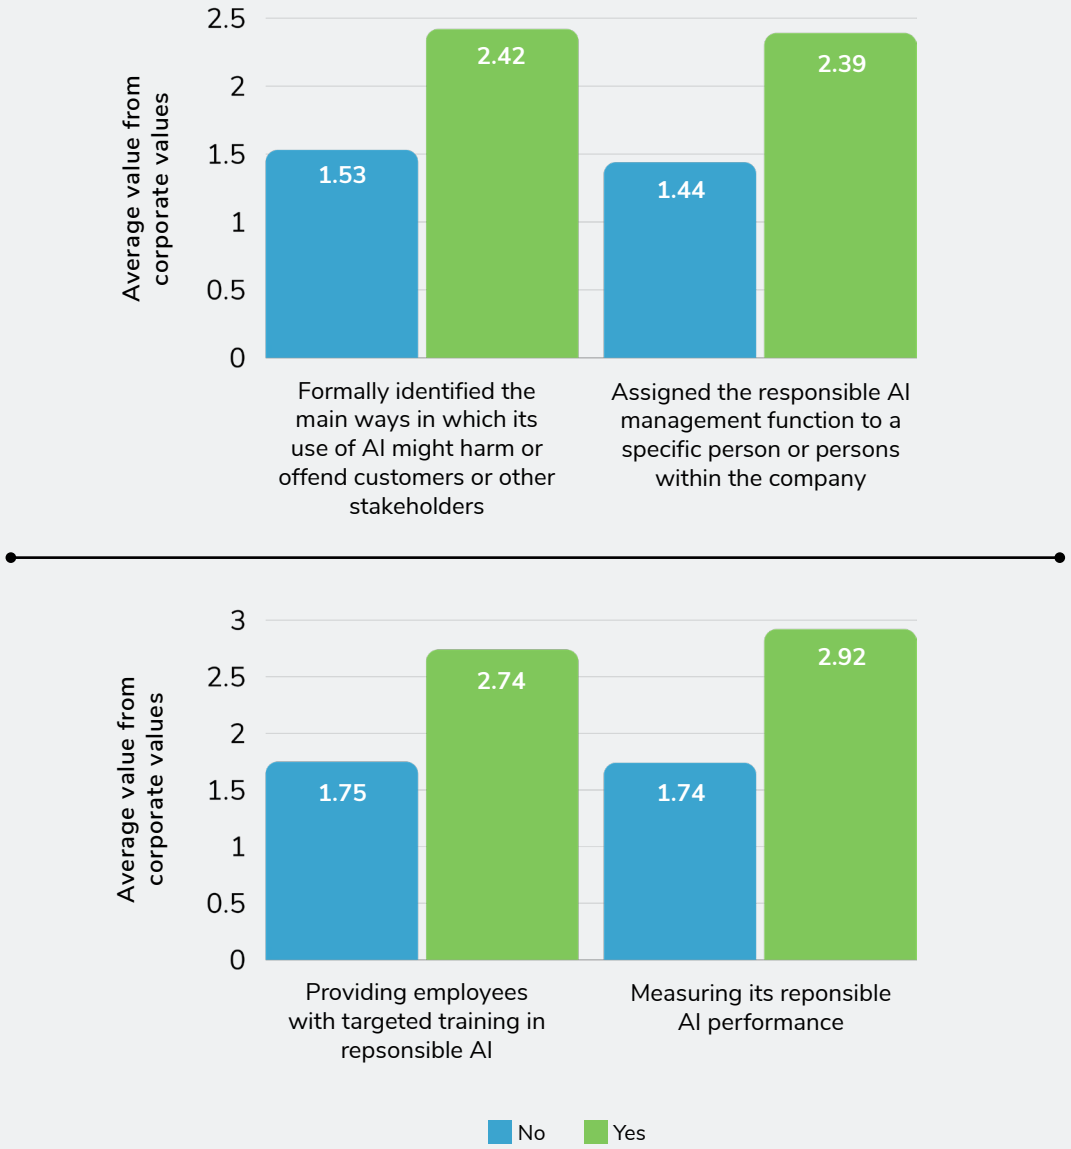

Considering the above relationships as a whole, it is worth noting that several of the RAIM activities that have the strongest relationship to particular types of value – requiring suppliers to follow the company's RAIM policies, training employees in RAIM, and measuring the company's RAIM performance – are also some of the least frequently implemented activities (Figure 8). Companies could potentially get more value from their RAIM programs by expanding them to include these activities.

The findings in this section only begin to explore the relationships between RAIM practices and RAIM value. Further research should be able to identify more such correlations and, of equal importance, start to reveal the mechanisms by which particular practices yield specific types of value. Such research will take us farther down the path towards identifying RAIM best practices, and so to research that can inform more fully organizational practice and public policy on responsible AI management.

CONCLUSION

This study set out to answer three questions:

What do businesses do when they pursue RAIM?

The survey results identify fourteen activities that are present in at least some business RAIM programs. They suggest that risk assessment, the building of a management structure, and the establishment of standards (ethical principles, policies) are the most common components of business RAIM initiatives. By contrast, measuring the firm's RAIM performance, evaluating employees' RAIM performance, and requiring a company's suppliers to follow its RAIM policies are among the least common.

Who in the business is responsible for RAIM?

The survey responses suggest that businesses are most likely to vest RAIM responsibility in those with expertise in privacy governance. That is perhaps unsurprising given the connections between privacy governance and AI governance. Firms also assign RAIM responsibility to those with data analytics, risk, and ethics expertise. This suggests that organizations may benefit from a multi-disciplinary approach to AI governance.

What value does RAIM create for the companies that implement it?

Respondents clearly stated that RAIM produced substantial value and that it did so in areas important to business strategy and competitiveness such as product quality, trustworthiness, and reducing regulatory risk. Businesses contemplating whether to implement their own RAIM program, and policymakers making the case for RAIM requirements, should find this relevant.

These results, while clearly preliminary, offer a tantalizing glimpse of a new way of thinking about responsible AI management that goes beyond the realm of corporate values and places it squarely at the center of an organization's AI strategy. They point the way towards future research that can more fully document, evaluate, and understand this important new governance area.

THE OHIO STATE UNIVERSITY

PROGRAM ON DATA AND GOVERNANCE
A program of the Moritz College of Law and the Translational Data Analytics Institute

iapp

Dennis D. Hirsch
Faculty Director, Program on Data and Governance
Core Faculty Member, Translational Data Analytics Institute
Hirsch.151@osu.edu

Angie Westover-Muñoz
Program Manager, Program on Data and Governance
westover-munoz.1@osu.edu

https://moritzlaw.osu.edu/PDG