*** DRAFT – PRELIMINARY AND INCOMPLETE***

Report: Self-governing ethical AI development in entrepreneurship

May 1, 2022

James Bessen, Boston University

Stephen Michael Impink, New York University

Robert Seamans, New York University

Abstract: This report summarizes key results from our fourth annual AI startups survey and analyzes a new survey question on governance-related practices that startups employ to address fairness and algorithmic bias. We find that AI startups with ethical AI policies are more likely to have corporate governance practices, including internal audits, external audits, AI oversight boards, guiding their AI product development. Next, we show that existing relationships with technology firms, notably Amazon, Google, and Microsoft, are related to increased corporate governance around AI for startups. Furthermore, we find ethical AI policies are related to increased product development-related governance practices (Human-in-the-loop, A/B testing) when startups have more advanced technological capabilities (neural networks, ensemble learning, AI frameworks). Lastly, we map these corporate and product development-related governance increases to startups dropping more biased training in development.

Introduction

Entrepreneurs developing AI confront numerous competing issues impacting their ventures' survival (Bessen et al., 2020). Startups must access the training data needed for product development while adhering to existing regulations (e.g., GDPR, CCPA), understanding pending regulations (e.g., AI governance in Europe), and rising to the ethical norms of their industry to maintain their legitimacy with potential investors. However, despite the well-documented risk of AI harming demographic subgroups (Barocas and boyd, 2017; Cowgill and Tucker, 2019; Friedman and Nissenbaum, 1996), particularly when training AI with personal data (Barocas and Nissenbaum, 2014; Barocas and Levy, 2020; boyd and Crawford, 2012; Whittaker et al., 2018), no widespread regulations exist.

Dealing with big data, algorithmic bias, and AI-based discrimination raises concerns about existing corporate governance models, which traditionally focus on maximizing shareholders' value (Cihon et al., 2021; Gibney, 2020). Ethical AI development principles (Bessen et al., 2022a; Mittelstadt, 2019; Schiff et al., 2020) and industry standards (Winfield, 2019) have emerged as mechanisms to help private firms and industries govern AI, but documentation may not be enough to safeguard consumer welfare (Mittelstadt 2019). Firms must also structure their organization to fit the constantly changing technological environment as AI advances (Eitel-Porter, 2021; Hilb, 2020; Lipai et al., 2021; Mäntymäki et al., 2022; Stahl et al., 2021).

This paper explores governance mechanisms that AI startups establish when developing their products. Our research question asks whether and how particular governance mechanisms available to startups are related to more ethical AI development outcomes. Ethical AI policies are a step toward developing more ethical products; however, all policies are likely different, and it is unlikely that policies are implemented uniformly across firms. This variance in policy content and implementation across firms is likely wider for startups than for more established firms. We survey 412 AI-producing startups to learn more about their ethical considerations, policies, and governance mechanisms. We then build measures of startups that implement corporate governance practices (e.g., internal audits, external audits, and AI oversight board) or product-development-related governance practices (e.g., A/B texting, human-in-the-loop) to reduce algorithmic bias and mitigate ethical issues in AI's production and use.

Survey Data and Summaries

We use data from a survey of AI startups, including questions regarding the development of ethical AI principles, the impact of those ethical principles, and actions and mechanisms the startups enact to be more ethical. We pretested the survey instrument with several academics and practitioners associated with startups and then administered the survey on two occasions: January 2021 and March 2022. We received

412 responses from AI startups in our sample; these firms confirmed they develop AI products in the first survey question.

Respondents to our survey come from several sampling frames, including Crunchbase, Pitchbook, Creative Destruction Labs, and an incubator at Technische Universität München (TUM). From Crunchbase and Pitchbook, we identify firms associated with the keyword "artificial intelligence" that are in operation yet have not experienced an initial public offering (IPO). Additionally, we received a contact list of AI startups from the Creative Destruction Lab, a startup incubator based in Toronto, and another contact list from Philipp Hartmann and Joachim Henkel at TUM (Hartmann and Henkel 2018). To develop a more homogenous sample of AI startup firms for our analysis, we exclude larger (more than 500 employees) and older (more than ten years old) firms.

We summarize descriptive statistics for startups included in surveys 3 and 4 in Table 1. In Figures 1-6, we report summaries of survey data related to startups' ethical development of AI. In the appendix, we report additional summaries for 1) combined data from all four rounds of the survey (A1-A10), 2) survey 3 and 4 data on exit plans (A11-A14), and 3) survey 4 data on AI frameworks (A15-16).

Methods

We use regression models to explore the relationship between ethical AI policies and governance. We use CEM weighting and matching (Iacus et al., 2019), based on region, startup age, employment size, and VC funding before the survey was administered, to ensure that the firms with an ethical AI policy are observationally similar to those without an ethical AI policy. The matching approach reduces the difference in standardized means across these observable demographic variables between the respondents who have and do not have ethical AI principles (Table 1, under matched)

We use the following regression specification for our analysis:

$$(1)\ y_i = \beta_0 + \beta_1 prior\_resource_i + \rho_i + \mu$$

where, y_i refers to an indicator variable for if the firm uses a governance mechanism (yes, 1; otherwise, 0); ai_policy_i refers to an indicator variable if a firm has an ethical AI policy; ρ are controls for small employment size (<11 employees) and HQ location in a top VC city; μ is the error term; we use robust standard errors in all regressions. We control for firm size and if a firm is located in a city with a high concentration of venture capital firms, which may be related to outcomes. For example, smaller firms have fewer employees, so they may be less likely to fire employees. Additionally, location in the same city with many venture capital firms may make it more likely that a startup can access capital or learn about good market opportunities.

Regression Results

We find that having an ethical AI policy is correlated with three corporate governance mechanisms: internal audits (Table 2 model (1): +0.29 SD 0.06), external audits (model (4): +0.11 SD 0.04), and having an AI oversight board (model (7): +0.16 SD 0.04). The relationship between having an ethical AI policy and these governance mechanisms increases if startups have a data-sharing or supplier relationship with a Big Tech firm (i.e., Amazon, Google, or Microsoft) (Table 3). Next, we analyze product-related governance practices and find that having an ethical AI policy is only correlated with product-related ethics practices when startups have a certain level of technical capacity, like using more sophisticated algorithms or AI development frameworks (including TensorFlow, PyTorch, etc.) (Table 4). Lastly, these governance mechanisms are positively correlated to dropping more biased data (Table 5).

These survey results provide insight into how AI startups, an important source of AI product innovation, manage ethical issues with little regulatory oversight while lacking the resources of established firms. Relationships with large technology firms are correlated with startups having an ethical AI policy (Bessen et al., 2022b) and implementing corporate governance mechanisms focused on reducing algorithmic bias. These findings spotlight the importance of norms in the IT services industry on the nascent AI industry and their interconnectivity. Next, the development of more sophisticated products, often requiring more data and technical know-how, creates a need for increased governance. Our results show a relationship between the increased use of advanced AI frameworks or more sophisticated algorithms and the increased use of product-related governance mechanisms, like human-in-the-loop and A/B testing.. Lastly, these results show a relationship between startups using these governance mechanisms and dropping more biased data, providing insight into the tough decisions these newly formed startups encounter, weighing the tradeoff between increased production and ethicality that materially impacts their business.

Figures and Tables

Figures 1 and 2: Prior ethics questions: These figures use combined data from survey 3 and 4.

Figure A.11 – Firm Actions by Ethics Policy Adoptions

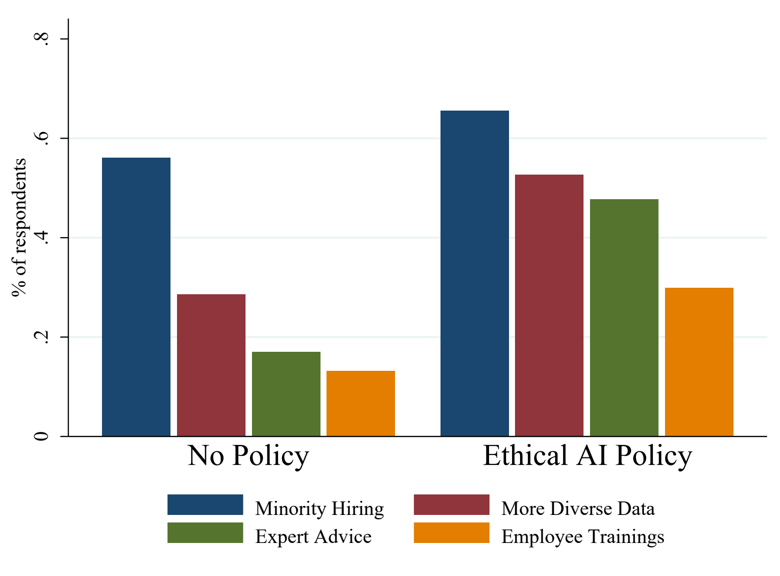

Figure A.11 – Ethical Outcomes

All startups included have an ethical AI policy

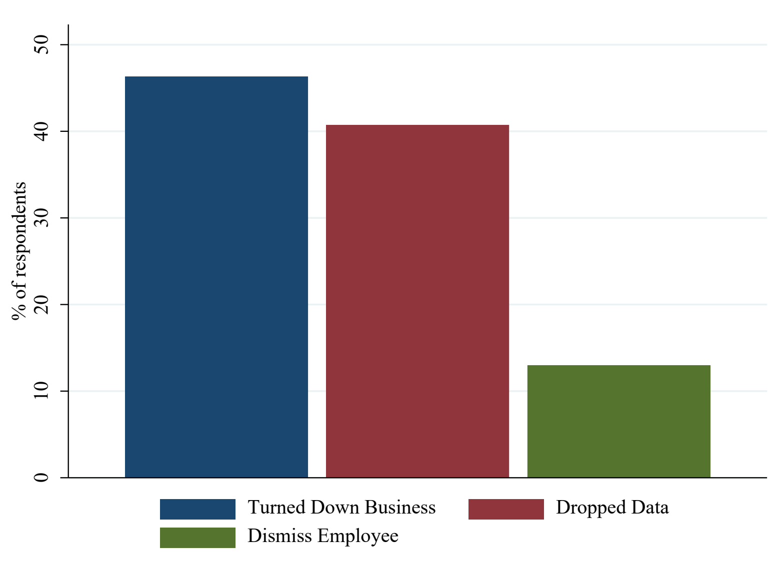

Figures 3-6: New ethics question: These figures use data from survey 4.

Figure 3 – Ethical Practices

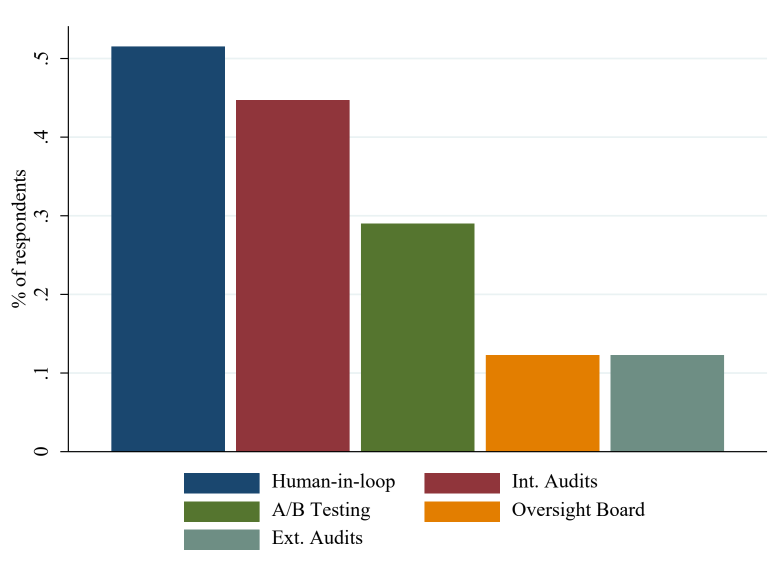

Figure 4 – Ethical Practices by Size

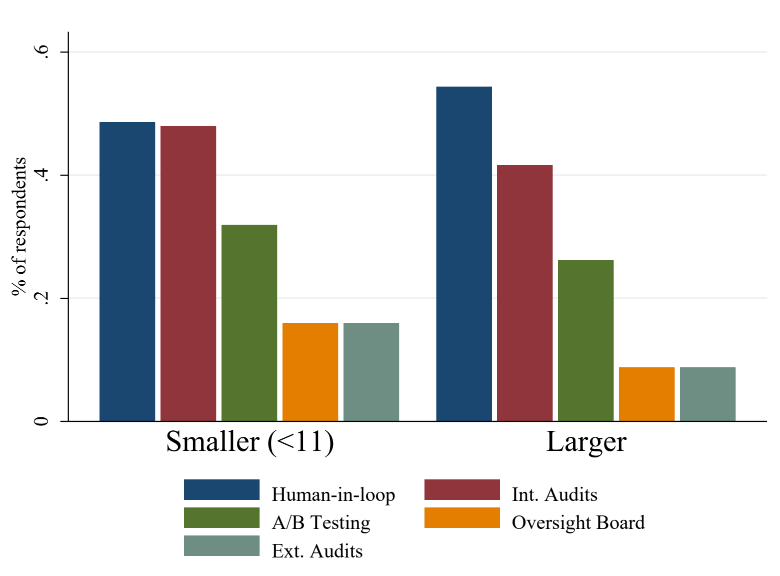

Figure 5 – Ethical Practices by Data Sharing Collaboration

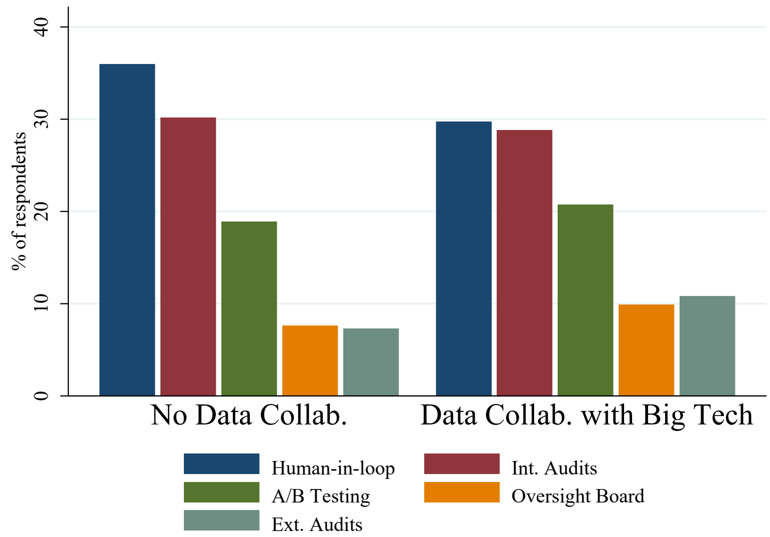

Figure 6 – Ethical Practices by Region

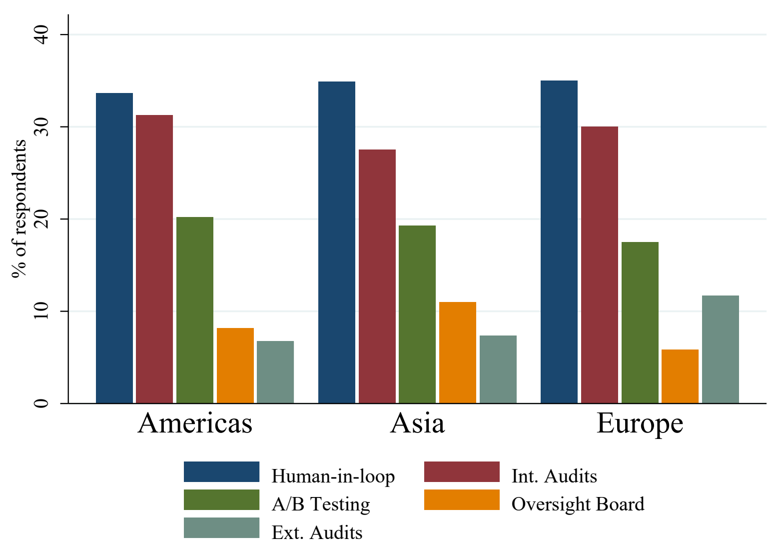

Table 1 – Descriptive Statistics

CEM	Unmatched								Matched			
Sample:	All				Policy		No Policy		Policy		No Policy	
	Mean	SD	Min	Max	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Age	3.78	2.19	0	9	4.01	2.32	3.61	2.08	3.74	2.17	3.6	2.08
Age (<3)	0.09	0.29	0	1	0.08	0.27	0.11	0.31	0.08	0.27	0.11	0.31
Employees	24.16	33.85	1	375	22.81	30.03	25.18	36.5	20.69	22.27	21.79	22.8
Employees (<11)	0.49	0.5	0	1	0.49	0.5	0.49	0.5	0.52	0.5	0.5	0.5
Financial	0.06	0.24	0	1	0.05	0.21	0.08	0.27	0.05	0.22	0.08	0.27
Healthcare	0.13	0.34	0	1	0.12	0.33	0.14	0.34	0.13	0.33	0.14	0.35
BT CSP	0.34	0.47	0	1	0.38	0.49	0.31	0.46	0.38	0.49	0.32	0.47
Funded	0.5	0.5	0	1	0.54	0.5	0.47	0.5	0.52	0.5	0.48	0.5
VC Backed	0.34	0.48	0	1	0.41	0.49	0.29	0.46	0.39	0.49	0.3	0.46
Revenue	0.39	0.49	0	1	0.45	0.5	0.34	0.48	0.43	0.5	0.35	0.48
BT Funding	0.02	0.14	0	1	0.02	0.15	0.02	0.13	0.02	0.15	0.02	0.13
US	0.41	0.49	0	1	0.42	0.5	0.4	0.49	0.42	0.5	0.4	0.49
Germany	0.04	0.19	0	1	0.03	0.18	0.04	0.2	0.04	0.19	0.04	0.2
France	0.02	0.15	0	1	0.02	0.15	0.03	0.16	0.02	0.15	0.03	0.16
UK	0.04	0.2	0	1	0.03	0.18	0.05	0.22	0.03	0.18	0.05	0.22
Canada	0.04	0.2	0	1	0.05	0.22	0.03	0.18	0.06	0.23	0.03	0.18
Americas	0.49	0.5	0	1	0.5	0.5	0.47	0.5	0.5	0.5	0.48	0.5
Asia	0.18	0.39	0	1	0.16	0.37	0.2	0.4	0.17	0.38	0.2	0.4
EU	0.31	0.46	0	1	0.31	0.46	0.31	0.47	0.31	0.47	0.31	0.46
MEA	0.02	0.13	0	1	0.02	0.15	0.01	0.11	0.01	0.09	0.01	0.09
California	0.13	0.34	0	1	0.16	0.37	0.11	0.31	0.16	0.37	0.11	0.32
New York	0.06	0.24	0	1	0.05	0.22	0.07	0.25	0.06	0.24	0.07	0.25
Massachusetts	0.02	0.14	0	1	0.01	0.11	0.03	0.16	0.01	0.1	0.03	0.16
San Francisco	0.07	0.26	0	1	0.1	0.3	0.05	0.22	0.09	0.29	0.05	0.22
N		412			177		235		176		230	

Table 2 - Corporate Governance and AI Policies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
DV, Dummy:	Internal Audits			External Audits			AI Oversight Board		
Ethical AI Policy	0.294***	0.301***	0.298***	0.105***	0.104**	0.111***	0.163***	0.159***	0.165***
	(0.061)	(0.062)	(0.061)	(0.040)	(0.041)	(0.040)	(0.040)	(0.041)	(0.041)
Tech Data		-0.041			0.010			0.020	
Sharing Coll.		(0.062)			(0.042)			(0.043)	
BT CSP			0.099			0.109*			0.042
			(0.074)			(0.058)			(0.053)
Employees (<11)	-0.035	-0.035	-0.021	-0.065	-0.065	-0.049	-0.064	-0.064	-0.058
	(0.061)	(0.061)	(0.062)	(0.041)	(0.041)	(0.041)	(0.042)	(0.042)	(0.042)
VC Location	-0.015	-0.013	-0.033	-0.005	-0.005	-0.025	-0.013	-0.014	-0.021
	(0.084)	(0.084)	(0.085)	(0.055)	(0.055)	(0.058)	(0.054)	(0.054)	(0.057)
Firms	256	256	256	256	256	256	256	256	256
Adj R2	0.0758	0.0737	0.0788	0.0225	0.0189	0.0367	0.0530	0.0501	0.0517
CEM Weighted	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

Table 3 - Corporate Governance and AI Policies with Tech Relationship Mechanism Policies

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
DV, Dummy:	Internal Audits			External Audits			AI Oversight Board		
Ethical AI Policy	-0.069	0.263***	0.301***	0.009	0.100**	0.092**	0.028	0.130***	0.115**
	(0.091)	(0.069)	(0.070)	(0.051)	(0.042)	(0.042)	(0.044)	(0.046)	(0.046)
Tech Data Sharing Collaboration	0.276***			0.103*			0.166***		
	(0.089)			(0.057)			(0.059)		
AI Pol. x Tech Collaboration	0.257***			0.114**			0.181***		
	(0.081)			(0.052)			(0.052)		
BT Data Sharing Collaboration		-0.002			0.065			-0.049*	
		(0.123)			(0.082)			(0.028)	
AI Pol. x BT Collaboration		0.394***			0.173**			0.238***	
		(0.094)			(0.080)			(0.083)	
Big Tech CSP			0.106			0.066			-0.071**
			(0.109)			(0.068)			(0.034)
AI Pol. x BT CSP			0.395***			0.238***			0.257***
			(0.101)			(0.087)			(0.090)
Employees (<11)	-0.034	-0.022	-0.021	-0.065	-0.060	-0.048	-0.064	-0.051	-0.054
	(0.061)	(0.061)	(0.062)	(0.041)	(0.042)	(0.041)	(0.043)	(0.042)	(0.042)
VC Location	-0.017	-0.015	-0.034	-0.006	-0.004	-0.018	-0.013	-0.014	-0.003
	(0.085)	(0.084)	(0.086)	(0.056)	(0.055)	(0.059)	(0.054)	(0.053)	(0.058)
Firms	256	256	256	256	256	256	256	256	256
Adj R2	0.070	0.075	0.075	0.015	0.022	0.035	0.046	0.064	0.065
CEM Weighted	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: * p<0.1, ** p<0.05, *** p<0.01.

Table 4 - Prior Technical Resouces and Product-related Ethics Practices

	(1)	(2)	(3)	(4)	(5)	(6)
DV, Dummy:		A/B Testing			Human in the Loop	
Ethical AI Policy	0.010	0.108	0.193**	0.093	0.290**	0.196*
	(0.060)	(0.113)	(0.098)	(0.063)	(0.129)	(0.116)
Neural. Net or Ensemble Lrn.		0.158*			0.212*	
		(0.093)			(0.109)	
AI Pol. x NN or EL		0.134			0.244**	
		(0.086)			(0.105)	
AI Framework			0.268***			0.187*
			(0.083)			(0.099)
AI Pol. x AI Framework			0.180**			0.219**
			(0.071)			(0.091)
Employees (<11)	-0.083	-0.070	-0.083	0.030	0.042	0.031
	(0.059)	(0.061)	(0.059)	(0.063)	(0.064)	(0.062)
VC Location	0.049	0.056	0.020	0.028	0.036	0.005
	(0.083)	(0.083)	(0.083)	(0.085)	(0.086)	(0.086)
Firms	256	256	256	256	256	256
Adj R2	-0.002	-0.000608	0.0226	-0.00195	0.00598	0.00455
CEM Weighted	No	Yes	Yes	No	Yes	Yes

Notes: * p<0.1, ** p<0.05, *** p<0.01.

Table 5 - Corporate Governance and Outcomes

	(1)	(2)	(3)	(4)	(5)
DV, Dummy:			Drop Data		
Internal Audits	0.278***				
	(0.047)				
External Audits		0.213**			
		(0.092)			
AI Oversight			0.286***		
Board			(0.090)		
AB Testing				0.139**	
				(0.057)	
Human-in-the-loop					0.056
					(0.050)
Employees (<11)	-0.097**	-0.093*	-0.088*	-0.095*	-0.109**
	(0.047)	(0.049)	(0.049)	(0.049)	(0.050)
VC Location	0.075	0.068	0.072	0.058	0.064
	(0.070)	(0.073)	(0.073)	(0.073)	(0.072)
Firms	256	256	256	256	256
Adj R2	0.133	0.0424	0.0713	0.0371	0.0153
CEM Weighted	Yes	Yes	Yes	Yes	Yes

Notes: * p<0.1, ** p<0.05, *** p<0.01.

References

Barocas, S., & Levy, K. (2020). Privacy dependencies. Wash. L. Rev., 95, 555.

Barocas, S., and boyd, D. Engaging the ethics of data science in practice. Communications of the ACM, 60(11), 23-25. 2017.

Barocas, S., and Nissenbaum, H. Big data's end run around anonymity and consent. Privacy, big data, and the public good: Frameworks for engagement, 1, 44-75. 2014

Bessen, J. E., Impink, S. M., Reichensperger, L., & Seamans, R. (2018). The Business of AI Startups. Boston Univ. School of Law, Law and Economics Research Paper, (18-28).

Bessen, J., Impink, S. M., Reichensperger, L., & Seamans, R. (2022a). The role of data for AI startup growth. Research Policy, 51(5), 104513.

Bessen, J., Impink, S.M. & Seamans, R. (2022b, forthcoming). The Cost of Ethical AI Development for Startups. In Proceedings of 2022 ACM Conference on Artificial Intelligence, Ethics, and Society (AIES'22), August 1–3, 2022, Oxford, United Kingdom.

boyd, D., and Crawford, K. Critical Questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon. Information, Communication and Society 15, no. 5: 662-679. 2012.

Cihon, P., Schuett, J., & Baum, S. D. (2021). Corporate governance of Artificial Intelligence in the public interest. Information, 12(7), 275.

Cowgill, B., and Tucker, C. E. Economics, Fairness and Algorithmic Bias. Journal of Economic Perspectives. 2019.

Eitel-Porter, R. (2021). Beyond the promise: implementing ethical AI. AI and Ethics, 1(1), 73-80.

Friedman, B., and Nissenbaum, H. Bias In Computer Systems. ACM Transactions on Information Systems (TOIS) 14, no. 3, 30-347. 1996.

Gibney, E. (2020). The battle for ethical AI at the world's biggest machine-learning conference. Nature, 577(7791), 609-610.

Iacus, S. M., King, G., and Porro, G. Causal inference without balance checking: Coarsened exact matching. Political Analysis, 1-24. 2012.

Lipai, Z., Xiqiang, X., & Mengyuan, L. (2021). Corporate governance reform in the era of artificial intelligence: research overview and prospects based on knowledge graph. Annals of Operations Research, 1-16.

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). Defining organizational AI governance. AI and Ethics, 1-7.

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., and Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2), 1-16.

Schiff, D., Biddle, J., Borenstein, J., & Laas, K. (2020, February). What's next for ai ethics, policy, and governance? a global overview. In Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (pp. 153-158).

Stahl, B. C. (2021). Artificial intelligence for a better future: an ecosystem perspective on the ethics of AI and emerging digital technologies (p. 124). Springer Nature.

Whittaker, M., Crawford, K., and Dobbe, R. Genevieve Fried, Elizabeth Kaziunas, Varoon Mathur, Sarah Mysers West, Rashida Richardson, Jason Schultz, and Oscar Schwartz. AI Now Report 2018. New York: AI Now Institute at New York University. 2018.

Winfield, A. (2019). Ethical standards in robotics and AI. Nature Electronics, 2(2), 46-48.

Appendix

Appendix A – Overview of all four rounds of the AI startup survey

Figures A.1-A.2: The survey has reached 765 unique startups out of a population of around 4,300 AI startups identified in our sample from Crunchbase and Pitchbook. About 18% of identified AI-producing startups have responded to at least one round of the survey. More than 200 firms have responded to more than one round of the survey.

Figure A.1 – Response count by survey

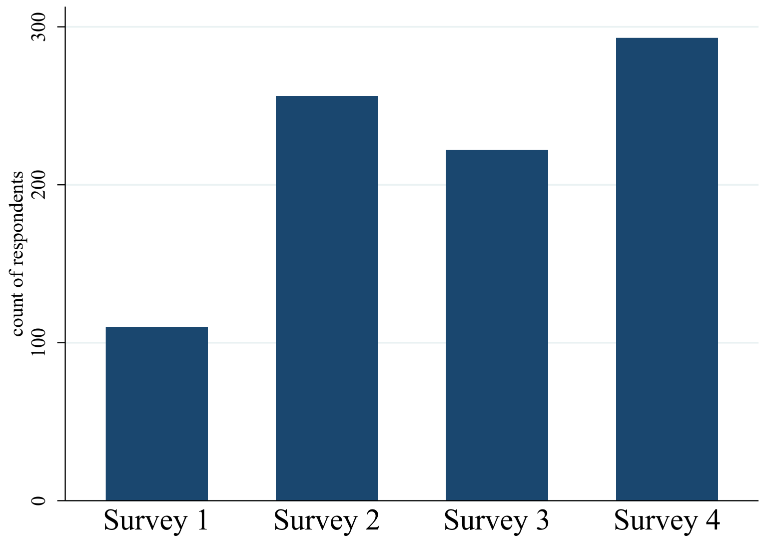

Figure A.2 – Repeat responders

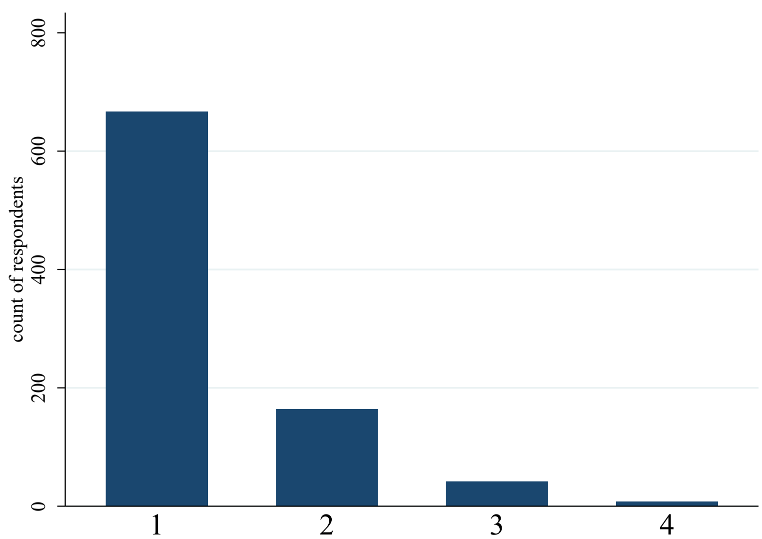

Figures A.3-A.10: The following views have been complied across all four rounds of the survey.

Figure A.3 – Product Status

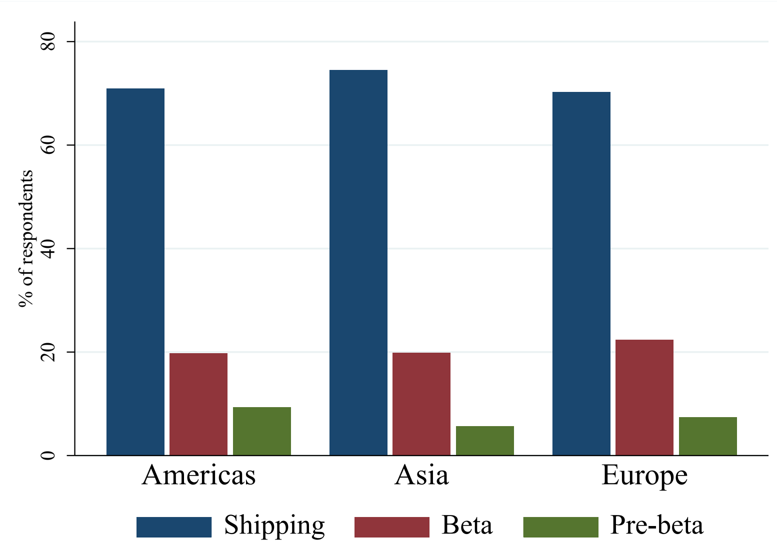

Figure A.4 – Customer Industries

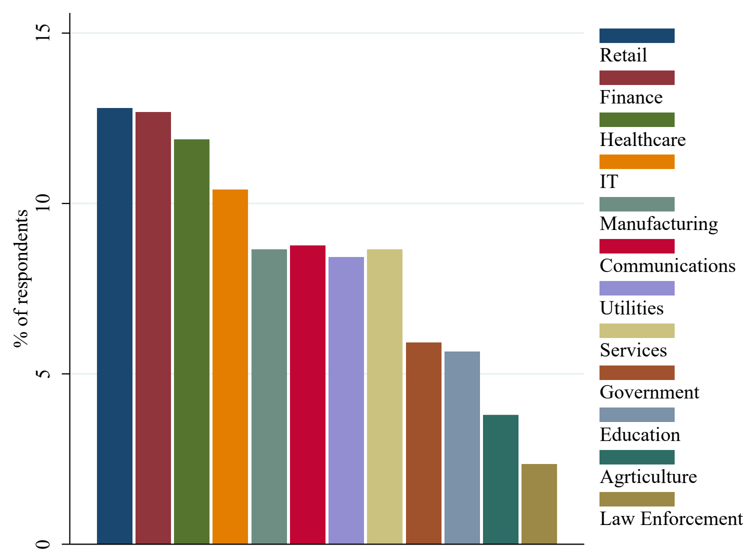

Figure A.5 – Customer Occupations

Who uses the AI product in your customers firms?

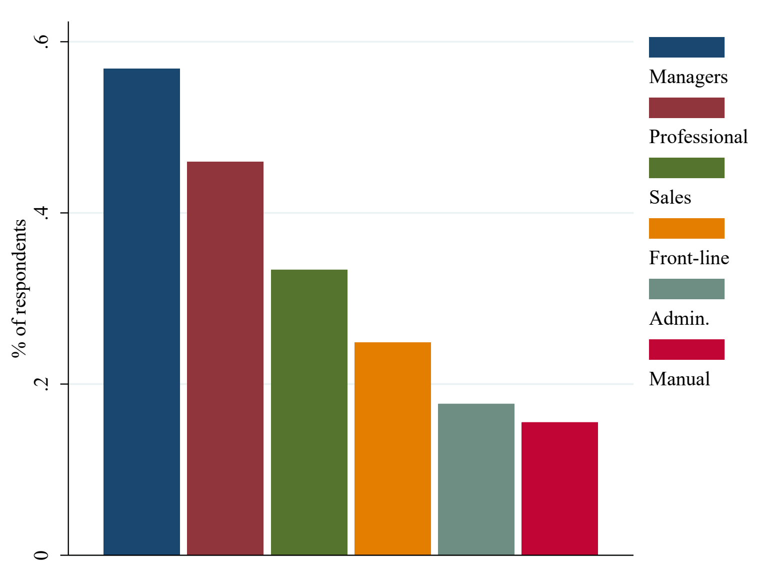

Figure A.6 – Headquarters Regions

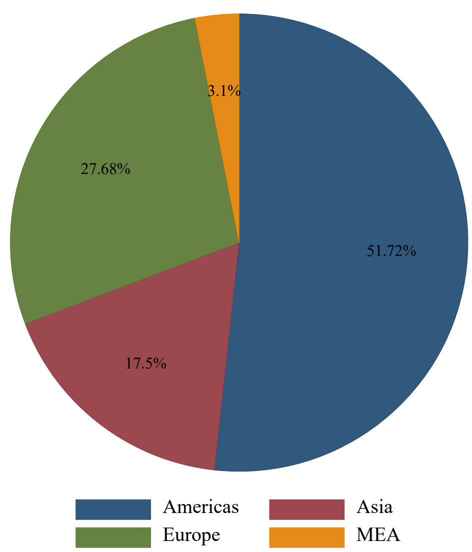

Figure A.7 – Age

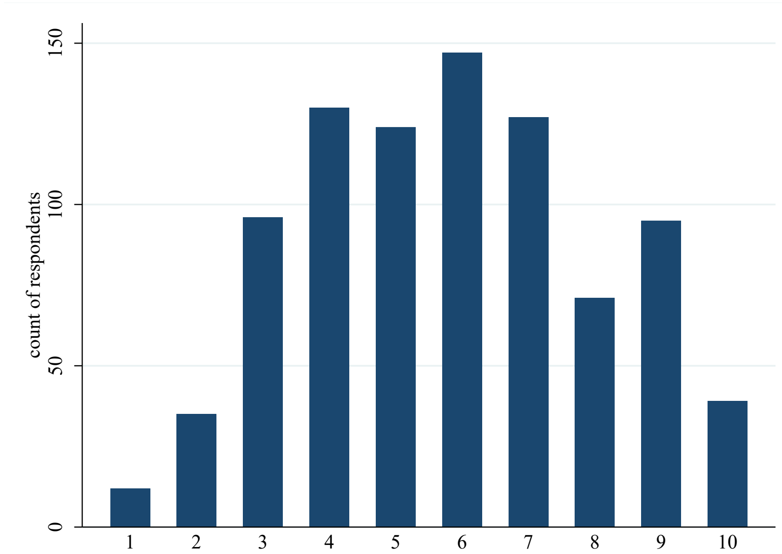

Figure A.8 – Cloud Usage by Provider

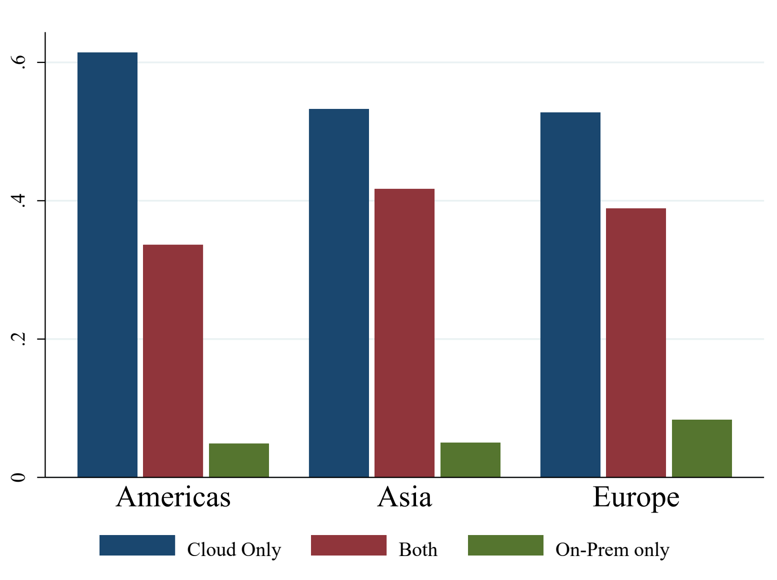

Figure A.9 – Cloud Provider by Size

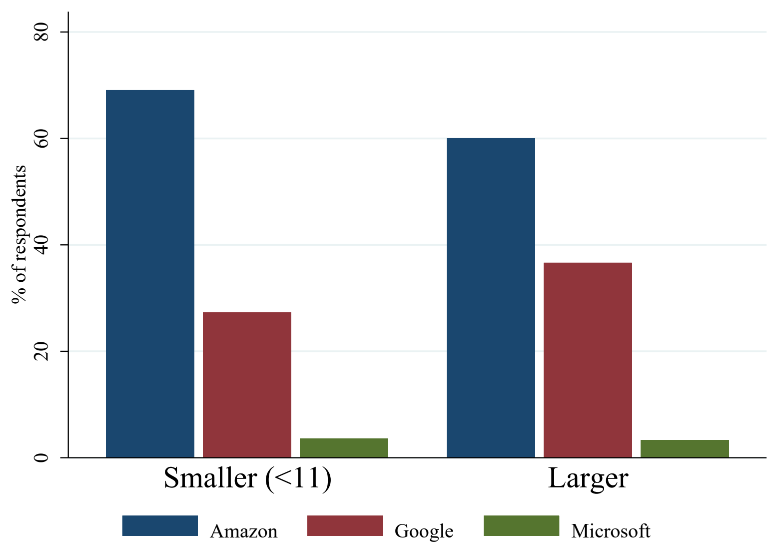

Figure A.10 – Cloud Provider by Region

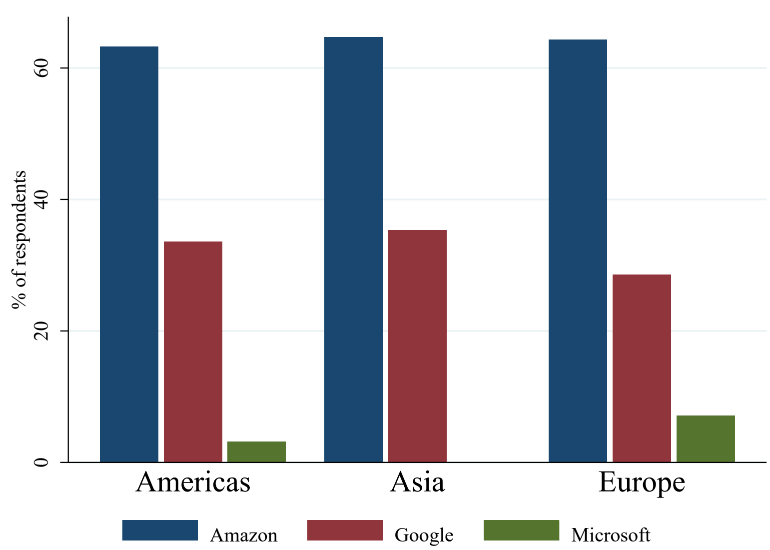

Figures A.11-A.14: Prior exit plan questions: These figures use combined data from survey 3 and 4.

Figure A.11 – Exit Plans by Fit

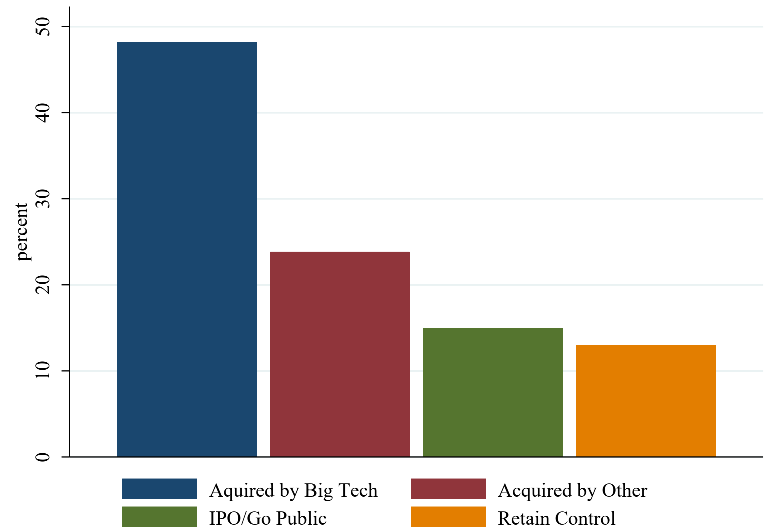

Figure A.12 – Exit Plans by Region

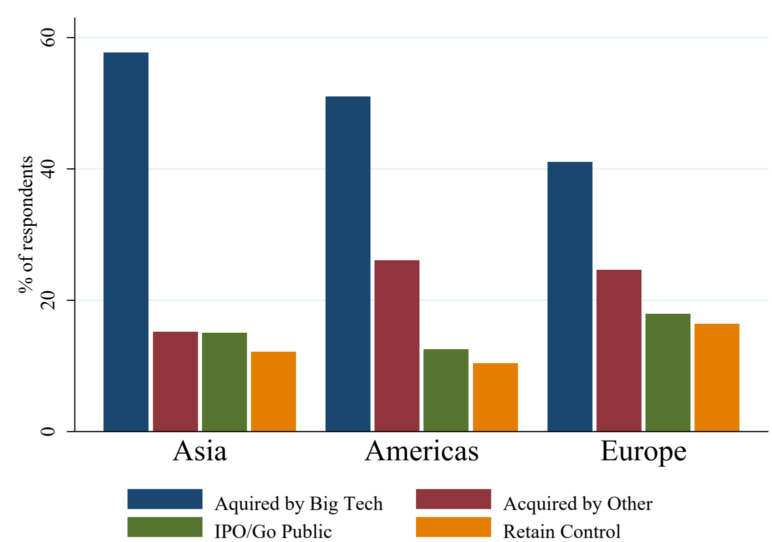

Figure A.13 – Exit Plans by Fit

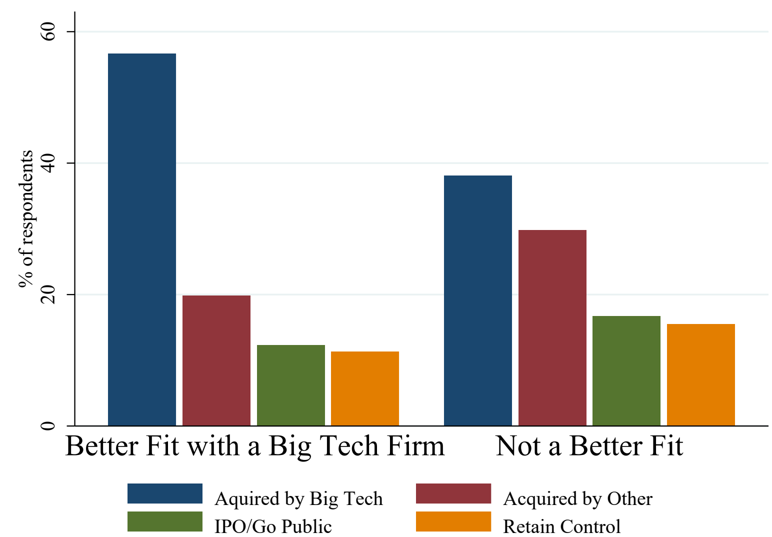

Figure A. 14 – Exit Plans by Data Sharing Collaboration

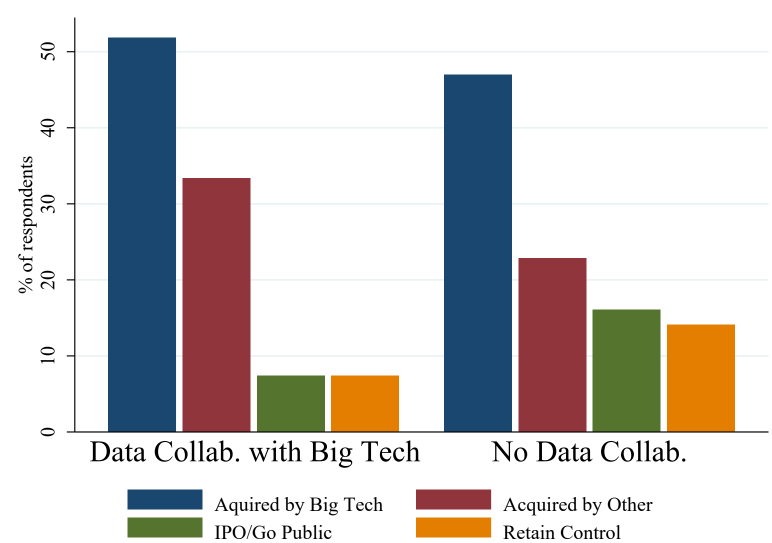

Figures A.15-A.16: New frameworks question: These figures use data from 4.

Figure A.15 – AI Development Frameworks

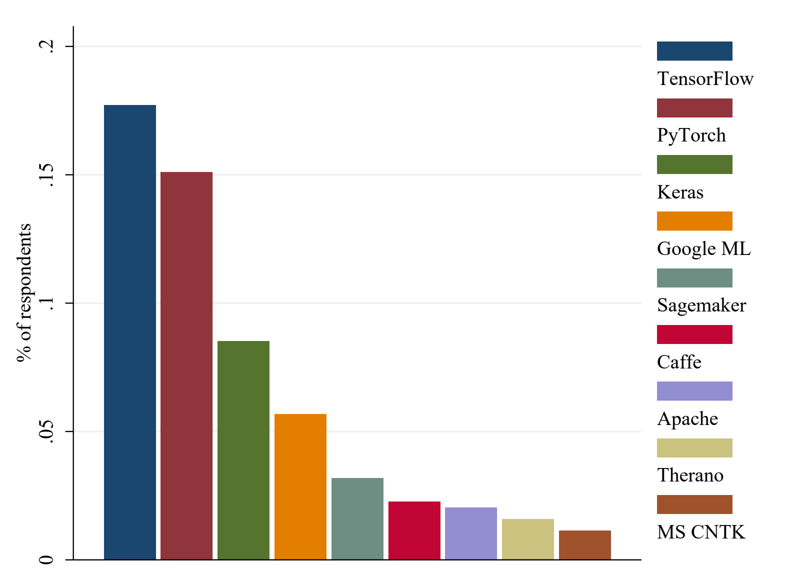

Figure A.16 – AI Framework by Algorithm Type