Post-Deployment Regulatory Oversight for General-Purpose Large Language Models

Carson Ezell
Harvard University
cezell@college.harvard.edu

Abraham Loeb
Harvard University
aloeb@cfa.harvard.edu

Abstract

The development and deployment of increasingly capable, general-purpose large language models (LLMs) has led to a wide array of risks and harms from automation that are correlated across sectors and use cases. Effective regulation and oversight of general-purpose AI (GPAI) requires the ability to monitor, investigate, and respond to risks and harms that appear across use cases, as well as hold upstream developers accountable for downstream harms that result from their decisions and practices. We argue that existing processes for sector-specific AI oversight in the U.S. should be complemented by post-deployment oversight to address risks and harms specifically from GPAI usage. We examine oversight processes implemented by other federal agencies as precedents for the GPAI oversight activities that a regulatory agency can conduct. The post-deployment oversight function of a regulatory agency can complement other GPAI-related regulatory functions that federal regulatory agencies may perform which are discussed elsewhere in the literature, including pre-deployment licensing or model evaluations for LLMs.

1 Introduction

The development of increasingly capable large language models (LLMs) trained on general-domain text corpora, which can be used or adapted for a wide array of tasks (48), has led to greater consolidation in the artificial intelligence (AI) development and usage landscape (49). In particular, rather than separate AI systems being developed for differing contexts without many shared components (e.g. datasets), similar or identical LLMs are increasingly used across contexts, and different systems often share components (24; 4; 5).

AI regulation in the U.S. often involves targeting particular applications through sector-specific regulatory agencies, rather than regulating general-purpose systems or the AI development process itself (23; 25). However, these recent trends in AI development have sparked significant regulatory discussion about how to ensure that upstream development of general-purpose systems is conducted responsibly (37); that users across applications are prepared to appropriately integrate general-purpose models (27); and that a unified, cross-sectoral approach to addressing the impacts of advanced AI is developed (38).

Experts tend to agree that regulation of general-purpose systems should include regulations on AI development and release as well as oversight of deployed systems in particular contexts (34). Some proposals for regulating development involve measures such as licensing model development and deployment (2), or conducting pre-deployment audits and evaluations (35). Conversely, proposals for oversight of deployed systems tend to rely more heavily on empowering sector-specific regulators to implement use case regulations (14).

This workshop paper argues that a federal regulatory agency should perform post-deployment oversight functions to address risks from general-purpose AI, in close collaboration with sector-

37th Conference on Neural Information Processing Systems (NeurIPS 2023).

specific regulators. Risks and harms discovered in particular applications may stem from aspects of the upstream development process, which suggests these failures would be correlated across usage contexts (3). Identifying these correlations would result in a more widespread and comprehensive understanding of GPAI-induced risks and harms. Furthermore, such an understanding would enable holding upstream developers of general-purpose models—who are usually well-resourced, private technology companies (1)—accountable for harms downstream of their development practices (36).

2 Proposal

We argue that a regulatory agency with the authority to address risks and harms from general-purpose LLM development and deployment could effectively carry out post-deployment monitoring and oversight of GPAI in coordination with sector-specific regulators. Some post-deployment responsibilities of an agency may include:

1. **Monitoring risks and harms:** Aggregating evidence of GPAI-related risks and harms to identify patterns and inform investigations or standards-setting related to the upstream development process.
2. **Facilitating investigation and corrective action:** Enforcing corrective action on general-purpose model developers or deployers when usage data reveals that their models create unacceptable risks and harms or fail to satisfy certain requirements or standards.
3. **Documenting interdependencies:** Improving traceability of harms throughout the model development and deployment process, including linking harms discovered in certain contexts to general-purpose model developers' practices or underlying model dependencies (5).
4. **Improving government accountability:** Ensuring that government agencies using GPAI systems and sector-specific regulators meet their transparency, oversight, and reporting obligations (25) by providing assistance and identifying shortcomings.
5. **Improving cross-sectoral coordination:** Facilitating information-sharing about risks and harms from general-purpose models across various sectors.
6. **Enabling public engagement:** Creating avenues for the public to inform oversight, investigations, corrective action, and rulemaking for GPAI based on observed risks and harms.

The oversight activities we propose are precedented by federal agencies that conduct oversight in other areas. To develop our proposals, we examined seven U.S. regulatory agencies that we classified across four categories which share some similarities to AI oversight, as shown in Table 1: consumer-product monitoring, critical infrastructure monitoring, market monitoring, and online monitoring.

Table 1: Categorization of regulatory agencies we analyzed to inform our proposals

Category	Agencies
Consumer Products	National Highway Traffic Safety Administration (NHTSA), Food and Drug Administration (FDA)
Critical Infrastructure	Federal Aviation Administration (FAA), Nuclear Regulatory Commission (NRC)
Markets	Federal Regulatory Energy Commission (FERC), Commodity Futures Trading Commission (CFTC)
Online	Cybersecurity and Infrastructure Security Agency (CISA)

We propose several related processes that a regulatory agency could implement based on the precedents set by other agencies: incident reporting, information sharing, petitions, investigations, and corrective action. For each of these processes, we (1) identify shared, standard procedures across federal agencies and (2) provide recommendations for how similar procedures can enable effective GPAI oversight. Figure 1 depicts the role of an agency conducting GPAI oversight within the broader post-deployment oversight landscape that we envision.

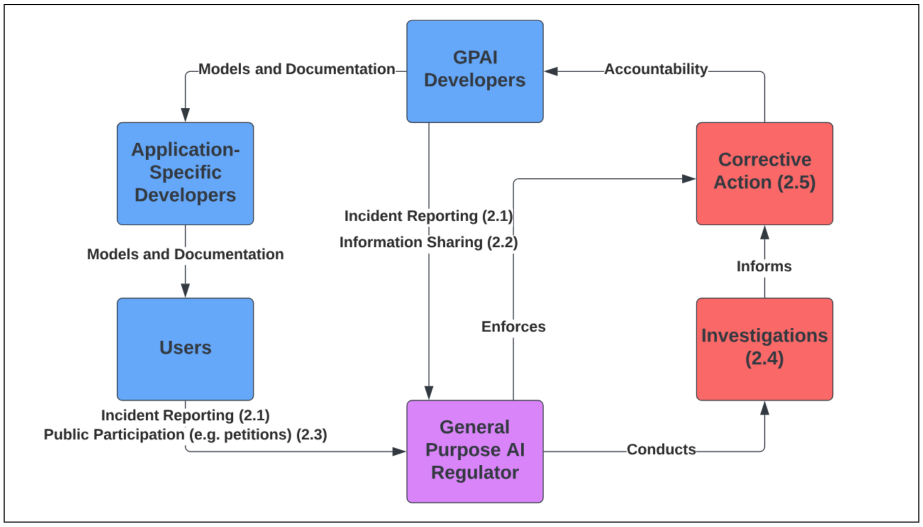

Figure 1: Relationship between GPAI regulator and key processes and stakeholders for post-deployment oversight. Numbers refer to subsections of this paper.

2.1 Incident Reporting

A regulatory agency could maintain a reporting system and database for general-purpose LLM-related incidents, with required incident reporting for LLM developers, deployers, high-stakes users, and government agencies. Each agency we examined has similar reporting requirements for risks, incidents, and/or standards violations which may apply to licensees (28; 16; 41; 6), manufacturers (47; 44), or government users (11). For example, the FDA requires manufacturers to report previously unknown interactions or adverse effects from medications, which are then shared through the FAERS database (22). NHTSA requires manufacturers to report vehicle complaints, communications with dealers and owners about defects, and other forms of information that might suggest the existence of defects (47).

Other government agencies or regulatory authorities could also be obligated to share information about LLM-related incidents that result from their own LLM usage. This would parallel the requirement that Federal Civilian Executive Branch (FCEB) agencies report cybersecurity incidents to CISA within an hour of their occurrence (11). Furthermore, regulatory bodies could be required to report LLM-related incidents that occur at entities they oversee—energy transmission system operators, known as Regional Transmission Organizations (RTOs), are required to have Market Monitoring Units (MMUs) that report observed violations of standards related to energy infrastructure or markets to FERC (17).

Employees of LLM developers, deployers, or high-risk users should also be protected when voluntarily raising concerns to regulators, similar to whistleblower protections for raising concerns to employers or regulators established by the NRC (28). Furthermore, penalties for self-reporting violations of standards could be met with reduced penalties to incentivize more reporting, based on policies that have been established at FERC (17) and the FAA (45).

Most agencies we examined also maintain voluntary reporting systems for public incident reporting (16; 21; 46; 32; 43; 13). Incident reports are usually shared publicly, often through periodic reports or an online database, with necessary privacy measures taken (46; 17; 30; 12; 26; 43). The CPSC maintains a platform for manufacturers and retailers to issue public replies to complaints (43)—a similar system could be used in the LLM context for upstream developers to provide explanations for their models' behavior in response to incident reports.

A regulatory agency could also develop recommendations and requirements for registering and tracking models to facilitate aggregation of reported risks and harms from the same models. Documentation and tracking of model dependencies (e.g. datasets) would enable greater traceability of harms and linkage between models that rely on the same dependencies (5), or are modified variants of the same foundation model. Such a system would resemble vehicle and vehicle component tracking by NHTSA (46), or product identification systems recommended by CPSC (9). All agencies we examined have some mechanisms to link reported risks and harms to their sources, including registration (6), licensing (29), product identification systems (9), or required identification of impacted systems (11).

2.2 Information Sharing

Developers, deployers, or high-risk users of LLMs can have obligations to share certain information with the AI Agency about regulated activities. This information may include statistics about model usage or details about risk management processes. These actors can also be required to store more sensitive information, such as training datasets and model inference logs, for extended periods of time (e.g. over a year) or indefinitely in case the information becomes relevant to an investigation and is requested by regulators.

Government agencies regularly obtain streams of information from regulated actors to assist oversight. This may include high-frequency or real-time information sharing about direct conduct of an overseen activity. For example, financial markets report daily positions and transactions for futures and options contracts for each clearing member to the CFTC (6). Similarly, scheduled electric power transactions in wholesale markets have electronic tags (eTags) which detail information about the transaction that is shared with FERC, and energy markets also provide FERC with daily information about market bids, offers, and outcomes (17). Through the Automated Indicator Sharing (AIS) program, data relevant to cybersecurity from participating entities is automatically shared with CISA and relevant partners (12).

High-risk stakeholders can provide more detailed reporting of their activities, including developers, auditors (10), or users who may have access to high-risk systems with potentially dangerous capabilities (35). Analogously, large options traders—defined as those who exceed certain reporting thresholds—are required to fill out more detailed forms to enable better tracking of their trading activities (7).

Other government agencies can also be required to report relevant information, including their own usage of AI systems (25) or known information about levels of automation or usage of LLMs in particular sectors. Measuring the extent of usage may be relevant for investigations into general-purpose systems, especially assessing the potential severity of risks and harms. These assessments would resemble how the FDA collects medication sales distribution data, hospital discharge billing data, insurance claims data, and health records which can give it a better sense of risk exposure to particular medications (44).

2.3 Petitions

Under the Administrative Procedure Act, agencies are required to allow members of the public to petition for a rule to be issued, amended or repealed (40). An AI Agency with sufficient investigative resources could facilitate public participation in GPAI rulemaking by having adequate resources to process petitions for rule updates based on observed risks or harms that result from inadequacies in existing rules.

A regulatory agency can also process Defect Petitions (42) which would call for particular corrective action by a model developer or deployer in response to observed model inadequacies. NHTSA allows members of the public to submit Defect Petitions, which call on NHTSA to open a defect investigation and provide a justification in the event of denial of a petition (42).

Furthermore, a regulatory agency could enable members of the public to submit petitions for enforcement for some action (e.g. civil penalty) to be taken against a developer or deployer. The NRC (29) and FDA (20) process petitions for enforcement action against licensed entities. When a petition is received by the NRC, a Petition Review Board may investigate the petition and issue a recommendation to relevant NRC leadership regarding whether enforcement action should be taken (29).

2.4 Investigations

A regulatory agency should have the sufficient talent, authority, and financial resources to gather necessary information and conduct a comprehensive investigation in the event of suspected wrongdoing on the part of upstream LLM developers or deployers. Most agencies conducting oversight have resources to conduct in-depth investigations of risks or regulatory violations within their scope, which often result from evidence gathered through incident reports or other information-sharing mechanisms (44; 46; 31; 13; 17). NHTSA has a dedicated Trends Analysis Division which reviews aggregate data and identifies defects for further investigation (47). Such a model may be appropriate for a regulatory agency that aggregates harms associated with general-purpose LLMs from many use cases.

Investigations often involve preliminary risk classifications to assess severity and urgency (9; 47), some degree of public transparency and reporting, and consideration of appropriate enforcement actions. When in-depth investigations are needed, investigators often conduct site visits (8), request additional information from relevant parties, and conduct rigorous tests (47). NHTSA investigations often involve Engineering Analysis in test centers, including frequent collaboration with third-parties (47). Similarly, third-parties in the AI auditing ecosystem may be well-positioned to collaborate with regulatory officials on investigations that involve rigorous model evaluations (33; 23).

A regulatory agency may also require upstream LLM developers to conduct further investigations themselves and release results if new potential risks are discovered. The FDA requires many studies and clinical trials of drug sponsors to investigate risks in the post-market phase, which may be required at the time of approval or upon discovery of new risks (39).

The agency can also conduct further post-deployment audits or investigations of upstream developers that are either random or targeted towards higher-risk or higher-impact systems. FERC conducts many audits based on perceived risk, even if there is no clear evidence for wrongdoing (17).

2.5 Corrective Action

Following the completion of an investigation, a regulatory agency could have an internal committee of experts decide the appropriate corrective action. Possible corrective action that may be required of upstream LLM developers includes improving documentation or communication of risks, conducting further model evaluations, improving safety guardrails or security practices, altering access protocols, withdrawing a model from usage, or deleting copies of model weights entirely.

Agencies often assemble committees of internal experts to determine what corrective action is necessary if an investigation identifies risks or wrongdoing (18; 45; 46; 31; 13). For example, the FDA forms an ad hoc committee of internal scientists when deciding whether to request, and how to classify, a recall (18). Similarly, the NRC uses Allegation Review Boards to investigate and make determinations about incidents (31), and NHTSA consults a panel of internal experts when deciding whether to concur with investigators about issuing a recall (46).

When corrective action is required, agencies often assist or approve corrective action plans for entities required to carry them out, and they often oversee the implementation of the plan through reporting requirements to ensure it is adequately completed (19; 9; 46; 15).

3 Conclusion

This paper argues that it would be both sensible and effective for a regulatory agency to conduct post-deployment oversight of general-purpose LLMs, which may complement other GPAI regulatory functions that the same or separate agencies may conduct which are discussed elsewhere in the literature (e.g. pre-deployment measures). However, many important questions need to be answered before effective processes for post-deployment GPAI oversight can be implemented. First, there are still many open questions about which regulatory agency would be well-positioned to conduct such oversight activities (or whether a new agency is needed), as well as its scope, authority, and design. Furthermore, there are a lack of regulations and standards for general-purpose LLM development, which would enable the agency to enforce compliance with best practices in addition to risk monitoring. Finally, identifying the sources of downstream harms from GPAI and the correlations between harms that cut across use cases is challenging. Greater attention from a wide variety of stakeholders—including upstream developers, downstream developers, and users—is needed to aggregate more information about harms and build a better understanding of their causes.

Acknowledgments and Disclosure of Funding

This work was partially supported by the Institute for Theory and Computation at Harvard University.

References

[1] Nur Ahmed, Muntasir Wahed, and Neil C. Thompson. The growing influence of industry in AI research. *Science*, March 2023. Publisher: American Association for the Advancement of Science. URL: https://www.science.org/doi/10.1126/science.ade2420, https://doi.org/10.1126/science.ade2420 doi:10.1126/science.ade2420.

[2] Markus Anderljung, Joslyn Barnhart, Anton Korinek, and Jade Leung et al. Frontier AI Regulation: Managing Emerging Risks to Public Safety, July 2023. arXiv:2307.03718 [cs]. URL: http://arxiv.org/abs/2307.03718, https://doi.org/10.48550/arXiv.2307.03718 doi:10.48550/arXiv.2307.03718.

[3] Rishi Bommasani, Kathleen A. Creel, Ananya Kumar, Dan Jurafsky, and Percy Liang. Picking on the Same Person: Does Algorithmic Monoculture lead to Outcome Homogenization?, November 2022. arXiv:2211.13972 [cs]. URL: http://arxiv.org/abs/2211.13972, https://doi.org/10.48550/arXiv.2211.13972 doi:10.48550/arXiv.2211.13972.

[4] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, and Russ Altman et al. On the Opportunities and Risks of Foundation Models, July 2022. arXiv:2108.07258 [cs]. URL: http://arxiv.org/abs/2108.07258, https://doi.org/10.48550/arXiv.2108.07258 doi:10.48550/arXiv.2108.07258.

[5] Rishi Bommasani, Dilara Soylu, Thomas I. Liao, Kathleen A. Creel, and Percy Liang. Ecosystem Graphs: The Social Footprint of Foundation Models, March 2023. arXiv:2303.15772 [cs]. URL: http://arxiv.org/abs/2303.15772, https://doi.org/10.48550/arXiv.2303.15772 doi:10.48550/arXiv.2303.15772.

[6] Commodity Futures Trading Commission. CFTC Market Surveillance Program, 2023. URL: https://www.cftc.gov/IndustryOversight/MarketSurveillance/CFTCMarketSurveillanceProgram/index.htm.

[7] Commodity Futures Trading Commission. Large Trader Reporting Program, 2023. URL: https://www.cftc.gov/IndustryOversight/MarketSurveillance/LargeTraderReportingProgram/ltrp.html.

[8] Consumer Product Safety Commission, Office of Compliance and Field Operations. The Regulated Products Handbook, May 2013. URL: https://www.cpsc.gov/s3fs-public/pdfs/blk_pdf_RegulatedProductsHandbook.pdf.

[9] Consumer Product Safety Commission, Office of Compliance and Field Operations. Product Safety Planning, Reporting, and Recall Handbook, September 2021. URL: https://www.cpsc.gov/s3fs-public/CPSCRecallHandbookSeptember2021.pdf.

[10] Sasha Costanza-Chock, Inioluwa Deborah Raji, and Joy Buolamwini. Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22, pages 1571–1583, New York, NY, USA, June 2022. Association for Computing Machinery. https://doi.org/10.1145/3531146.3533213 doi:10.1145/3531146.3533213.

[11] Cybersecurity and Infrastructure Security Agency. US-CERT Federal Incident Notification Guidelines, April 2017. URL: https://www.cisa.gov/federal-incident-notification-guidelines.

[12] Cybersecurity and Infrastructure Security Agency. Cybersecurity Incident and Vulnerability Response Playbooks, November 2021. URL: https://www.cisa.gov/sites/default/files/publications/Federal_Government_Cybersecurity_Incident_and_Vulnerability_Response_Playbooks_508C.pdf.

[13] Department of Homeland Security, Cybersecurity and Infrastructure Security Agency. National Cyber Incident Response Plan, December 2016. URL: https://www.cisa.gov/sites/default/files/ncirp/National_Cyber_Incident_Response_Plan.pdf.

[14] Alex Engler. A comprehensive and distributed approach to AI regulation, August 2023. URL: https://www.brookings.edu/articles/a-comprehensive-and-distributed-approach-to-ai-regulation/.

[15] Federal Aviation Administration. VDRP User Guide Release 5.0. URL: https://vdrp.faa.gov/UserGuide.pdf.

[16] Federal Aviation Administration. Mandatory and Voluntary Incident Reporting, February 2023. URL: https://www.faa.gov/hazmat/incident-reporting.

[17] Federal Energy Regulatory Commission, Office of Enforcement. 2022 Report on Enforcement. Docket No. AD07-13-016, November 2022. URL: https://www.ferc.gov/media/fy2022-oe-annual-report.

[18] Food and Drug Administration. 21 CFR 10.30, June 2023. URL: https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?FR=10.30.

[19] Food and Drug Administration. 21 CFR 7.41, June 2023. URL: https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?FR=7.41.

[20] Food and Drug Administration. 21 CFR 7.42, June 2023. URL: https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm?FR=7.42.

[21] Food and Drug Administration, Center for Drug Evaluation and Research. Questions and Answers on FDA's Adverse Event Reporting System (FAERS), May 2019. URL: https://www.fda.gov/drugs/surveillance/questions-and-answers-fdas-adverse-event-reporting-system-faers.

[22] Food and Drug Administration, Center for Drug Evaluation and Research. FDA Adverse Event Reporting System (FAERS) Public Dashboard, October 2021. URL: https://www.fda.gov/drugs/questions-and-answers-fdas-adverse-event-reporting-system-faers/fda-adverse-event-reporting-system-faers-public-dashboard.

[23] Gillian K. Hadfield and Jack Clark. Regulatory Markets: The Future of AI Governance, April 2023. arXiv:2304.04914 [cs, econ, q-fin]. URL: http://arxiv.org/abs/2304.04914, https://doi.org/10.48550/arXiv.2304.04914 doi:10.48550/arXiv.2304.04914.

[24] Bernard Koch, Emily Denton, Alex Hanna, and Jacob G. Foster. Reduced, Reused and Recycled: The Life of a Dataset in Machine Learning Research, December 2021. arXiv:2112.01716 [cs, stat]. URL: http://arxiv.org/abs/2112.01716, https://doi.org/10.48550/arXiv.2112.01716 doi:10.48550/arXiv.2112.01716.

[25] Christie Lawrence, Isaac Cui, and Daniel Ho. The Bureaucratic Challenge to AI Governance: An Empirical Assessment of Implementation at U.S. Federal Agencies. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, pages 606–652, New York, NY, USA, August 2023. Association for Computing Machinery. URL: https://dl.acm.org/doi/10.1145/3600211.3604701, https://doi.org/10.1145/3600211.3604701 doi:10.1145/3600211.3604701.

[26] NASA Aviation Safety Reporting System. HAZMAT Safety Reporting, 2023. URL: https://asrs.arc.nasa.gov/hazmat.html.

[27] National Institute for Standards and Technology. Biden-Harris Administration Announces New NIST Public Working Group on AI, June 2023. Last Modified: 2023-06-22T14:59-04:00. URL: https://www.nist.gov/news-events/news/2023/06/biden-harris-administration-announces-new-nist-public-working-group-ai.

[28] Nuclear Regulatory Commission. Reporting Safety Concerns to the NRC. NUREG BR-0240, August 2017. URL: https://www.nrc.gov/docs/ML1720/ML17208A272.pdf.

[29] Nuclear Regulatory Commission. Enforcement Petition Process, March 2019. URL: https://www.nrc.gov/docs/ML1907/ML19070A037.pdf.

[30] Nuclear Regulatory Commission. Allegations Annual Reports, July 2023. URL: https://www.nrc.gov/about-nrc/regulatory/allegations/guidedocs.html.

[31] Nuclear Regulatory Commission. Backgrounder on Allegations, March 2023. URL: https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/allegation-process-bg.html.

[32] Nuclear Regulatory Commission. Report a Safety or Security Concern, April 2023. URL: https://www.nrc.gov/about-nrc/regulatory/allegations/safety-concern.html.

[33] Inioluwa Deborah Raji, Peggy Xu, Colleen Honigsberg, and Daniel E. Ho. Outsider Oversight: Designing a Third Party Audit Ecosystem for AI Governance, June 2022. arXiv:2206.04737 [cs]. URL: http://arxiv.org/abs/2206.04737, https://doi.org/10.48550/arXiv.2206.04737 doi:10.48550/arXiv.2206.04737.

[34] Jonas Schuett, Noemi Dreksler, Markus Anderljung, and David McCaffary et al. Towards best practices in AGI safety and governance: A survey of expert opinion, May 2023. arXiv:2305.07153 [cs]. URL: http://arxiv.org/abs/2305.07153.

[35] Toby Shevlane, Sebastian Farquhar, Ben Garfinkel, and Mary Phuong et al. Model evaluation for extreme risks, May 2023. arXiv:2305.15324 [cs]. URL: http://arxiv.org/abs/2305.15324.

[36] Irene Solaiman, Zeerak Talat, William Agnew, Lama Ahmad, Dylan Baker, Su Lin Blodgett, Hal Daumé III, Jesse Dodge, Ellie Evans, Sara Hooker, Yacine Jernite, Alexandra Sasha Luccioni, Alberto Lusoli, Margaret Mitchell, Jessica Newman, Marie-Therese Png, Andrew Strait, and Apostol Vassilev. Evaluating the Social Impact of Generative AI Systems in Systems and Society, June 2023. arXiv:2306.05949 [cs]. URL: http://arxiv.org/abs/2306.05949, https://doi.org/10.48550/arXiv.2306.05949 doi:10.48550/arXiv.2306.05949.

[37] The White House. FACT SHEET: Biden-Harris Administration Secures Voluntary Commitments from Leading Artificial Intelligence Companies to Manage the Risks Posed by AI, July 2023. URL: https://www.whitehouse.gov/briefing-room/statements-releases/2023/07/21/fact-sheet-biden-harris-administration-secures-voluntary-commitments-from-leading-artificial-intelligence-companies-to-manage-the-risks-posed-by-ai/.

[38] The White House. FACT SHEET: Biden-Harris Administration Takes New Steps to Advance Responsible Artificial Intelligence Research, Development, and Deployment, May 2023. URL: https://www.whitehouse.gov/briefing-room/statements-releases/2023/05/23/fact-sheet-biden-harris-administration-takes-new-steps-to-advance-responsible-artificial-intelligence-research-development-and-deployment/.

[39] United States, Congress, House. 21 USC 355, Office of the Law Revision Counsel, August 2023. URL: uscode.house.gov.

[40] United States, Congress, House. 5 USC 553, Office of the Law Revision Counsel, August 2023. URL: https://uscode.house.gov/.

[41] United States, Department of Energy, Federal Energy Regulatory Commission. 18 CFR 12.10, Legal Information Institute, Cornell Law School, 2023. URL: https://www.law.cornell.edu/cfr/text/18/12.10.

[42] United States, Department of Transportation, National Highway Traffic Safety Administration. 49 CFR Part 552, National Archives, August 2023. URL: https://www.ecfr.gov/current/title-49/subtitle-B/chapter-V/part-552.

[43] U.S. Consumer Product Safety Commission. Implementation of a Searchable Consumer Product Safety Incident Database, September 2009. URL: https://www.cpsc.gov/s3fs-public/pdfs/blk_media_cpsia212.pdf.

[44] U.S. Department of Health and Human Services, Food and Drug Administration, Center for Drug Evaluation and Research, and Center for Biologics Evaluation and Research. Best Practices in Drug and Biological Product Postmarket Safety Surveillance for FDA Staff, November 2019. URL: https://www.fda.gov/media/130216/download.

[45] U.S. Department of Transportation, Federal Aviation Administration. Aviation Safety Action Program. Advisory Circular 120-66C, March 2020. URL: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_120-66C_(Edit).pdf.

[46] U.S. Department of Transportation, National Highway Traffic Safety Administration. Motor Vehicle Safety Defects and Recalls. DOT HS 808 795, August 2017. URL: https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/14218-mvsdefectsandrecalls_041619-v2-tag.pdf.

[47] U.S. Department of Transportation, National Highway Traffic Safety Administration. Risk-Based Processes for Safety Defect Analysis and Management of Recalls. DOT HS 812 984, November 2020. URL: https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/14895_odi_defectsrecallspubdoc_110520-v6a-tag.pdf.

[48] Jason Wei, Yi Tay, Rishi Bommasani, and Colin Raffel et al. Emergent Abilities of Large Language Models, October 2022. arXiv:2206.07682 [cs]. URL: http://arxiv.org/abs/2206.07682, https://doi.org/10.48550/arXiv.2206.07682 doi:10.48550/arXiv.2206.07682.

[49] Meredith Whittaker. The Steep Cost of Capture, 2021. URL: https://papers.ssrn.com/abstract=4135581.