Export Controls and Open-Weight AI:

A Framework for Weights, Outputs, and Downstream Conduct

By Joe Khawam

August 17, 2026

Comments on this draft paper are invited and can be sent to joe@lawreforminstitute.org.

Executive Summary

Export controls can serve a limited but important purpose in addressing risks from open-weight AI. The public release of general-purpose model weights can place them outside of item-based controls under the Export Administration Regulations (EAR). However, the publication of weights is just the first of several points in the open-weight lifecycle at which export controls may apply. The disclosure of model outputs and other downstream activity may also raise export-control issues. Export controls may reduce open-weight risks differently at each layer. This paper argues that the application of controls at each layer should be guided by a marginal-risk framework, under which export controls have the most value when the transfer or activity adds capability or access beyond the realistic alternatives and a reachable actor exists.

The paper begins by explaining the relevant technology and then analyzes how existing export-control law applies at each stage of the open-weight lifecycle. It explains the marginal-risk policy framework and how current market dynamics, particularly the availability of capable Chinese open models, limit the effectiveness of unilateral controls on U.S. weight releases as a means of denying capability or access. It then applies that legal and policy analysis to each layer. It concludes that, under current conditions, controls on the publication of general-purpose model weights would do little to reduce marginal risk and the EAR's "published" exclusion should remain available for such releases.

The paper further cautions against using existing export-control authorities to impose mandatory restraints on publication, given the statutory and constitutional vulnerabilities. If U.S. open models begin to provide threat-relevant capabilities beyond foreign alternatives, Congress should create a purpose-built authority before an urgent need arises, with clear standards, appropriate limits, and prompt judicial review. The export-control agencies should instead focus their use of existing authorities on high-consequence output disclosures and on downstream services that materially assist weapons-of-mass-destruction or military-intelligence activities or that create demonstrated risks to U.S. systems. Risks beyond the practical reach of export controls require other domestic authorities and sustained international cooperation.

I. Introduction

In June 2026, the Commerce Department informed Anthropic that licenses were required for exports, reexports, or in-country transfers of Claude Mythos 5 and Fable 5 to foreign persons.[1] Because Anthropic could not verify each user's nationality in real time, it disabled both models worldwide. The directive to Anthropic was issued ten days after Executive Order 14409, which called for voluntary pre-release review of certain frontier models while disclaiming mandatory licensing.[2] Although the order does not expressly distinguish between open- and

[1] Letter from Howard W. Lutnick, Sec'y of Com., to Dario Amodei, Chief Exec. Officer, Anthropic (June 12, 2026) (on file with author); Anthropic, *Statement on the U.S. Government Directive to Suspend Access to Fable 5 and Mythos 5* (June 12, 2026), https://www.anthropic.com/news/fable-mythos-access.

[2] Exec. Order No. 14,409, Promoting Advanced Artificial Intelligence Innovation and Security, § 3(a)–(c), 91 Fed. Reg. 34,565, 34,566 (June 5, 2026).

DRAFT

closed-weight models, reports indicate that the implementation framework developed by the Trump administration currently excludes open-weight models but that may be temporary.[3]

A system of voluntary pre-release review, backed by the possibility of an export-control directive that could require an already-released model to be pulled back from public access, could have practical force with a small group of frontier developers whose models remain under centralized control. However, such an approach would be harder to apply to the open-weight ecosystem. Repeated directives would amount to case-by-case improvisation, and each would raise statutory and First Amendment concerns. In addition, most of the government's practical leverage disappears after model weights are published. Once weights are released, an export-control directive cannot recall downloaded copies or restore centralized control. The loss of leverage would matter less if open models trailed far behind the frontier, but open-weight models are now estimated to trail the frontier in cyber capabilities by only four to seven months, and highly capable open-weight releases continue, particularly from China.[4] Washington and Beijing are reportedly considering new restrictions as they face growing pressure to address the security risks of open models.[5] Yet the availability of comparable foreign models can undermine any unilateral action.

This paper examines where U.S. export controls can reduce the risks associated with open-weight AI. It analyzes weights, outputs, and downstream conduct separately and, for each layer, assesses when existing controls apply and whether their benefits justify their costs. It concludes that controls are most useful where they reach an actor who can prevent or limit a transfer or activity that would provide the recipient with greater capabilities than the recipient could obtain from realistic alternatives—most notably, foreign open models. At the layer of model weights, that condition most likely applies before publication of a U.S. model that is materially ahead of foreign alternatives. At the output and conduct layers, subsequent actors may boost a recipient's capability at the margin by disclosing outputs or providing downstream services. Such activities can be constrained even if the model weights themselves are not controlled or originate outside the United States.

On the other hand, existing law does not provide a clean basis for categorically prohibiting the publication or download of model weights, and no authorities can realistically restore control over weights that have already been widely copied. If future developments create

DRAFT

---

[3] *See* Courtney Rozen, *Trump Advisers Tell AI Firms They Will Not Safety-Test Open-Weight Models, Sources Say*, Reuters (Aug. 4, 2026), https://www.reuters.com/business/trump-advisers-tell-ai-firms-they-will-not-safety-test-open-weight-models-2026-08-04/; Hugo Lowell, *The White House Is Going to Expand Its AI Policy*, Wired (Aug. 12, 2026), https://www.wired.com/story/the-white-house-is-going-to-expand-its-ai-policy/.

[4] AI Sec. Inst., *How Far Behind the Frontier Are Leading Open Weight Models on Cyber?* (July 17, 2026), https://www.aisi.gov.uk/blog/how-far-behind-the-frontier-are-leading-open-weight-models-on-cyber; Laurie Chen, *China's Moonshot Unveils World's Largest Open AI Model, Closing In on U.S. Rivals*, Reuters (July 17, 2026), https://www.reuters.com/world/china/chinas-moonshot-unveils-worlds-largest-open-ai-model-closing-us-rivals-2026-07-17/.

[5] Maria Curi, *The Secret Trump Administration Battle to Fight Chinese AI*, Axios (July 20, 2026), https://www.axios.com/2026/07/20/ai-us-china-open-source-kimi; Fanny Potkin, *Beijing Is Looking at Curbing Overseas Access to China's Top AI Models, Sources Say*, Reuters (July 7, 2026, updated July 9, 2026), https://www.reuters.com/world/beijing-is-looking-curbing-overseas-access-chinas-top-ai-models-sources-say-2026-07-07/.

a need to restrain publication of U.S. model weights, Congress should pass purpose-built legislation rather than expecting the executive branch to rely on current export-control laws.

Thus, given the distinct roles and limits of existing authorities, export-control policy—as well as licensing and enforcement efforts—should focus on the conduct most likely to affect the underlying national security goals. To analyze these nuances, this paper first addresses the technology, then existing law, and finally policy. Part II explains what open weights are and why public release of the weights limits the government's ability to exercise control. Parts III and IV examine how existing legal regimes apply across the model lifecycle—e.g., how publication of weights affects item-based controls, what conduct-based authorities are relevant, and what constitutional limits apply—and how export-control responsibility evolves at each stage. Part V then evaluates the extent to which export controls are likely to work in practice, focusing on how the marginal risk of open-weight models is affected by the availability of foreign models as capabilities evolve. Finally, Part VI applies that analysis separately to weights, outputs, and downstream conduct to identify where existing authorities can materially reduce risks and where purpose-built legislation, other domestic authorities, and international cooperation are needed.

II. What Are Open-Weight Models?

AI model weights are numerical parameters adjusted during training. The weights help determine how a model converts inputs into outputs.[6] Some developers keep those weights private and offer access only through a web interface or application programming interface (API), while others release them publicly as "open" weights.[7] Open-weight models allow downstream users to customize a pretrained model without bearing the significant cost of training one from scratch, thereby lowering the technical and computational barriers to building capable systems.[8] This approach allows scientists, researchers, artists, and other innovators greater access to AI tools for socially beneficial initiatives and entrepreneurial activities.[9]

The public release of model weights transfers control of the models to downstream users. Once weights are widely available, the original developer generally cannot recall them or decommission systems built from them.[10] Users can run an open-weight model on local hardware without sending data to the developer, a feature that is valuable for privacy-sensitive applications. But because usage data does not flow back to the developer, neither the developer nor the government can monitor use or intervene to address misuse.[11] Downstream fine-tuning

---

[6] Yann LeCun, Yoshua Bengio & Geoffrey Hinton, *Deep Learning*, 521 Nature 436, 436 (2015).
[7] Sayash Kapoor et al., *On the Societal Impact of Open Foundation Models*, AI as Normal Technology (Feb. 2024), https://www.normaltech.ai/p/on-the-societal-impact-of-open-foundation. "Open weight" is distinguishable from "open source." Open weight refers to making the model's trained parameters available, while open source generally involves a release of the code, weights, and information about training data under open terms.
[8] Edward J. Hu et al., *LoRA: Low-Rank Adaptation of Large Language Models*, Int'l Conf. on Learning Representations 1–2 (ICLR 2022), https://arxiv.org/pdf/2106.09685v2.
[9] Nat'l Telecomms. & Info. Admin., U.S. Dep't of Com., *Dual-Use Foundation Models with Widely Available Model Weights* 28–30 (July 2024), https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf (hereinafter NTIA Report).
[10] Madhulika Srikumar et al., *Risk Mitigation Strategies for the Open Foundation Model Value Chain* 18–19 (July 2024), https://partnershiponai.org/wp-content/uploads/dlm_uploads/2024/07/open-foundation-model-risk-mitigation_rev3-1.pdf.
[11] NTIA Report, *supra* note 9, at 14, 31.

DRAFT

can also weaken safeguards built into the models, even when the user does not intend to do so.[12] Open-weight models therefore create both societal benefits and security risks that closed models do not present in the same form.

III. Export Control Framework for Weights, Outputs, and Conduct

The EAR, which generally covers dual-use items, and the ITAR, which regulates defense articles and defense services, are the principal U.S. export-control regimes. This Part addresses how these and related authorities apply to weights, outputs, and other downstream conduct, how publication changes the analysis, and what statutory and constitutional limits exist.[13]

a. The Export Administration Regulations (EAR)

i. Model Weights as "Technology"

The EAR, administered by the Commerce Department's Bureau of Industry and Security (BIS), governs exports, reexports, and in-country transfers of commodities, software, and technology within its jurisdiction.[14] The EAR's rules on software and technology are both relevant for understanding how the regime applies to a deployed AI model. The inference software necessary to use a model contains executable instructions, and thus fits the EAR definition of software, which covers programs executable by a computer.[15] By contrast, the weights are recorded numerical parameters that the software loads and applies. Because the weights are numerical data rather than programs, they do not fit the definition of software; they may instead fit the definition of technology—i.e., information necessary for the development, production, use, operation, or other specified function of an item—here, the operation and use of the deployed model.[16] And, as discussed below, BIS has expressly treated model parameters covered by ECCN 4E091 as technology.[17]

ii. The "Published" Exclusion

Although model weights likely qualify as technology under the EAR, their publication may remove them from export controls. Section 734.7 provides that technology or software is "published"—and thus not subject to the EAR—when it has been made available to the public without restrictions on its further dissemination, including through posting on a publicly available internet site.[18] Publicly accessible weights ordinarily satisfy that standard if they are

---

[12] Xiangyu Qi et al., *Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To!*, Int'l Conf. on Learning Representations 3 (ICLR 2024), https://arxiv.org/pdf/2310.03693.

[13] For a detailed discussion of how export controls apply to AI outputs from closed-weight models, *see* Tim Schnabel & Joe Khawam, *AI Outputs and National Security Controls* 5–6 (Law Reform Inst., Oct. 14, 2025), https://lawreforminstitute.org/report101425.pdf [hereinafter LRI Report].

[14] 15 C.F.R. §§ 734.2(a), 734.3(a), 736.2(a)–(b).

[15] 15 C.F.R. § 772.1 (definitions of "program" and "software").

[16] *Id.* (definition of "technology").

[17] *See infra* Part III.a.iii; *see also* ECCN 4E091, 15 C.F.R. pt. 774, supp. no. 1; 15 C.F.R. § 738.2(b) (describing the structure of the Commerce Control List, in which the letter following the category number designates the product group, with "D" entries controlling software and "E" entries controlling technology). While it is reasonable to conclude that 4E091 has clarified BIS's view that weights are "technology" under the EAR, the entry does not authoritatively address the classification of weights falling outside its scope.

[18] 15 C.F.R. § 734.7(a); *see also id.* § 734.3(b)(3)(i).

DRAFT

released without restrictions. However, the license under which model weights are released may impose relevant restrictions on further dissemination.

Models released under standard open-source licenses (such as MIT, Apache 2.0, and GPL) can be redistributed without limits on classes of recipients, and thus present the clearest case for qualifying for the exclusion.[19] In contrast, models with more restrictive licenses present closer questions.[20] A condition that restricts use of the weights but permits their redistribution may be consistent with § 734.7. Likewise, license terms addressing attribution and notice are likely also consistent with the exclusion. On the other hand, a license term that limits who may copy or redistribute the weights is more difficult to reconcile with § 734.7. In any event, eligibility for the exclusion depends on the actual terms of the license and the public's ability to access the weights, not on whether a developer labels a model as open.

iii. Evolving Controls and the AI Diffusion Rule

In January 2025, near the end of the Biden administration, BIS issued the AI Diffusion Rule, an interim final rule that created the first Commerce Control List (CCL) entry explicitly covering AI model parameters.[21] The rule added ECCN 4E091, which covers the parameters of models trained with 10^26 or more computational operations, and it imposed a worldwide license requirement for the export of unpublished parameters within that ECCN.[22] However, after President Trump returned to office, Commerce announced in May 2025 that BIS would not enforce the rule and would rescind it, although the entry has not actually been removed from the CCL.[23] But regardless of the rule's status, the underlying analysis remains relevant.

DRAFT

---

[19] Open Source Initiative, *The Open Source Definition*, Req. No. 5 ("No Discrimination Against Persons or Groups"), https://opensource.org/osd (stating that to qualify as open source, a license "must not discriminate against any person or group of persons"); Open Source Initiative, *OSI Approved Licenses*, https://opensource.org/licenses (listing MIT, Apache 2.0, and GPL as approved licenses that conform to the Open Source Definition).

[20] *See, e.g.,* Meta Platforms, Inc., *Llama 4 Community License Agreement*, https://www.llama.com/llama4/license/ (conditioning redistribution on recipient acceptance of the Agreement and compliance with attribution requirements); Google LLC, *Gemma Terms of Use*, https://ai.google.dev/gemma/terms (last modified Apr. 1, 2026) (stating that the terms apply to earlier Gemma models and directing users to a separate Gemma 4 license); Nia Castelly et al., *Gemma 4: Expanding the Gemmaverse with Apache 2.0*, Google Open Source Blog (Apr. 2, 2026), https://opensource.googleblog.com/2026/03/gemma-4-expanding-the-gemmaverse-with-apache-20.html.

[21] Framework for Artificial Intelligence Diffusion, 90 Fed. Reg. 4,544 (Jan. 15, 2025) [hereinafter AI Diffusion Rule] (interim final rule).

[22] *Id.* at 4,555 (adding new ECCN 4E091).

[23] Press Release, Bureau of Indus. & Sec., U.S. Dep't of Com., Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls (May 13, 2025), https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens. As of July 2026, Commerce has not completed the intended rulemaking, and 4E091 remains on the Commerce Control List. *See* ECCN 4E091, 15 C.F.R. pt. 774, supp. no. 1. The Government Accountability Office later concluded that the blanket non-enforcement announcement was a rule for purposes of the Congressional Review Act, while recognizing that the planned rescission had not been completed. U.S. Dep't of Com., Bureau of Indus. & Sec.—Applicability of the Congressional Review Act to the Rescission of the Artificial Intelligence Diffusion Rule, B-337935 (Comp. Gen. May 12, 2026), https://www.gao.gov/products/b-337935. GAO distinguished between Commerce's announced non-enforcement policy, which it found to be a Congressional Review Act-covered rule, and Commerce's planned rescission of the rule, which GAO found was not yet final agency action ripe for CRA review because Commerce had not completed the rulemaking process.

In this framework, BIS treated publication of model weights as a key threshold, excluding from ECCN 4E091 any parameters that qualify as "published."[24] BIS explained its empirical premise—i.e., that no known published model had, at that point, been trained with more than 10^26 computational operations—as well as its policy premise—i.e., that the benefits of open publication outweighed the risks "for now." However, the rule excluded all parameters that qualify as "published" regardless of whether either premise remained accurate. A later model crossing the size threshold would still qualify for the exclusion even if the risk/benefit analysis had changed.[25] The rule also excluded closed models less powerful than the strongest published model, reflecting a judgment that export controls do not achieve their goals when comparable or superior capabilities are already publicly available.

iv. U.S. Person and Other Catch-All Controls

Part 744 of the EAR imposes "catch-all" license requirements that do not depend on how an item is classified; instead, they are triggered by the end use, end user, or destination of a transaction.[26] Under the Export Control Reform Act (ECRA), these catch-all requirements can take two different forms: controls on items subject to the EAR, and controls on the activities of U.S. persons.[27] Item-based catch-all controls apply only to items subject to the EAR, so they do not restrict open-weight models that qualify for the "published" exclusion.

Catch-all requirements on U.S.-person activities, by contrast, apply regardless of whether any particular item—such as model weights—is subject to the EAR. ECRA provides a closed list of end uses and end users that, under § 744.6 of the EAR, U.S. persons cannot "support" without a license. BIS cannot add new categories to this list, which includes specified nuclear, missile, chemical- or biological-weapons end uses; whole plants for chemical-weapons-precursors; and military-intelligence end uses and end users.[28]

Within those categories, BIS has discretion to define which activities qualify as "support." The definition currently encompasses activities such as transferring items not subject to the EAR that are located abroad, or performing services, with knowledge that the activity will assist or benefit a covered end use or end user.[29] In the advanced-semiconductor controls, BIS used § 744.6 to impose categorical license requirements rather than require case-specific proof that a U.S. person knows an activity supports a particular WMD program. BIS expressly linked advanced semiconductors to WMD concerns, reasoning that advanced computing and

DRAFT

---

[24] AI Diffusion Rule, *supra* note 21, at 4,579 (Note 1 to ECCN 4E091).
[25] *Id.* at 4,555.
[26] 15 C.F.R. pt. 744; *see generally* Emily S. Weinstein & Kevin Wolf, *For Export Controls on AI, Don't Forget the "Catch-All" Basics*, Ctr. for Sec. & Emerging Tech. (July 5, 2023), https://cset.georgetown.edu/article/dont-forget-the-catch-all-basics-ai-export-controls.
[27] 50 U.S.C. § 4812(a)(1)–(2).
[28] 50 U.S.C. § 4812(a)(2)(A)–(F), (b)(2) (listing the same categories implemented by BIS in § 744.6, except that BIS has not implemented U.S.-person controls for "foreign maritime nuclear projects" in § 4812(a)(2)(E)); *see also id.* § 4813(d); 15 C.F.R. §§ 734.5(a), 736.2(b)(7), 744.6(a)–(b).
[29] 15 C.F.R. § 744.6(b)(6).

semiconductor capacity can support WMD development and that China's military-civil fusion policies make specific end uses difficult to identify.[30]

Section 744.6 thus operates in two ways. First, some obligations arise case by case, when a U.S. person knows its activity may support a covered end use. Second, BIS defines certain covered activities in advance and requires licenses categorically, as it did in the semiconductor controls. These catch-all prohibitions are most likely to be relevant to open-weight models when a U.S. person provides assistance to a covered program or actor. A U.S. person, for example, may need a license to fine-tune a model for biological-weapons development or to host a model for a covered program.[31] BIS has specifically advised that U.S. persons may need authorization to support AI-model training for WMD or military-intelligence activities.[32] An even more difficult question is once the weights encode WMD-relevant capability past a particular threshold, whether their release by U.S. persons would constitute support and require a license. Because these controls regulate U.S. persons' conduct, they apply regardless of whether the model weights come from any particular country or are subject to the EAR. Part VI further considers how these controls could extend to services involving open weights.

v. The June 2026 Anthropic Letter

The June 2026 Anthropic letter illustrates the item-based specific-notice route for catch-all controls. Among other authorities, the letter invoked the EAR's military-intelligence end-use controls (§ 744.22).[33] Section 744.22(b) permits BIS to inform a party that a license is required for an export, reexport, or in-country transfer of an item subject to the EAR because of an unacceptable risk of military-intelligence use or diversion.[34] The Secretary of Commerce informed Anthropic that licenses were required for exports, reexports, transfers, deemed exports, and deemed reexports of its Claude Mythos 5 and Fable 5 models to all destinations and foreign persons.[35] Anthropic disabled worldwide access to these models after receiving the letter because it could not verify the nationality of every user in real time. Following negotiations with the U.S. government, Commerce lifted the restrictions later that month, and Anthropic restored access to the models.[36]

This incident has two primary implications for open-weight models. First, it shows that Commerce is willing to use individualized notice—an "is-informed" letter—to impose restrictions on an AI model whose weights remained unpublished. BIS could take the same approach, delivering an is-informed letter to an open-weight developer, to restrict weights that

---

[30] Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification, 87 Fed. Reg. 62,186, 62,187–88, 62,199–200 (Oct. 13, 2022); 15 C.F.R. § 744.6(c)(2)–(3).
[31] *See* 15 C.F.R. § 744.6(b)(3), (5)–(6).
[32] Bureau of Indus. & Sec., U.S. Dep't of Com., *BIS Policy Statement on Controls That May Apply to Advanced Computing Integrated Circuits and Other Commodities Used to Train AI Models* 1–2 (May 13, 2025), https://www.bis.gov/media/documents/ai-policy-statement-training-ai-models-may-13-2025.
[33] *See* Letter from Howard W. Lutnick, *supra* note 1.
[34] 15 C.F.R. § 744.22(b).
[35] Letter from Howard W. Lutnick, *supra* note 1 (citing 50 U.S.C. §§ 4813(a)(15), 4817(b)(1); 15 C.F.R. § 744.22(b)).
[36] Anthropic, *Redeploying Fable 5* (June 30, 2026), https://www.anthropic.com/news/redeploying-fable-5.

DRAFT

are intended for later open release—though it could do so only before publication, since § 734.7 would place published weights beyond the reach of the item-based controls the letter invoked.

Second, the letter's practical effect was to condition foreign persons' access to a hosted model on a license. On its face, the letter applied only to the transmission or release of "the model" and said nothing about hosted access. However, the letter did not clarify what was meant by "the model" and thus what EAR-controlled item was being exported or disclosed. Commerce may have used "the model" as an umbrella term to cover both (a) transfers or releases of the underlying model weights, which present a more conventional item-based issue because ECCN 4E091 treats covered weights as technology, and (b) foreign persons' access to the model through Anthropic's hosted service, which provides users with outputs without giving them access to the underlying weights or software and thus presents a more difficult legal question. The context of the letter strongly suggests that Commerce sought to reach the second activity in addition to the first. Anthropic disabled worldwide access to both the Mythos 5 and the Fable 5 models after a reported exchange with Commerce Secretary Howard Lutnick,[37] who apparently confirmed that the license requirement covered use of the hosted service by foreign persons.

That position is in tension with earlier BIS advisory opinions treating the provision of cloud-computing capacity as a service rather than an export, and concluding, in the hosted-software context, that remote use without download is not an export of the software.[38] Commentators have rightly questioned whether foreign access to the hosted models involved an export or release of any item subject to the EAR.[39] The letter therefore creates uncertainty as to whether BIS will treat foreign access to a hosted model as an export or release even when the provider does not transfer or disclose the underlying weights or software—a theory that could affect providers that offer hosted access to open-weight models. The better reading of the EAR's item-based provisions is that hosted access without a transfer of weights or software is a service, not an export or release of "the model." The letter was likely outcome-driven, seeking the worldwide suspension of the Mythos 5 and Fable 5 models without attempting to explain how that result could be reconciled with the EAR's definitions or BIS's prior guidance. Nonetheless, it will serve as an enforcement signal for U.S. providers of hosted model access. Until BIS clarifies its position, those providers must account for the risk that BIS will apply the same theory to their provision of services to foreign users.

DRAFT

---

[37] *See A Big Shift in the AI Race*, The AI Daily Brief (June 17, 2026), https://aidailybrief.beehiiv.com/p/a-big-shift-in-the-ai-race ("During his Friday night call, Dario Amodei reportedly told Lutnick 'This means we can't have the model out.' Lutnick's reply: 'That's the point.'").

[38] Bureau of Indus. & Sec., U.S. Dep't of Com., Advisory Opinion, Application of the Export Administration Regulations to Grid and Cloud Computing Services (Jan. 13, 2009) (concluding that the provision of computational capacity through grid or cloud computing is a service rather than an export by the provider); Bureau of Indus. & Sec., U.S. Dep't of Com., Advisory Opinion, Cloud Computing and Deemed Exports (Jan. 11, 2011) (concluding that a cloud computing provider need not obtain deemed-export licenses for foreign-national IT administrators that monitor or screen user-generated technology); Bureau of Indus. & Sec., U.S. Dep't of Com., Advisory Opinion, Cloud-Based Storefronts (Nov. 13, 2014) (concluding that a foreign user's remote use of hosted software, without downloading it, is not an export of the software).

[39] *See* Kate Koren et al., *The Department of Commerce Restricted Access to Anthropic's Latest Models. What Comes Next?*, CSIS (June 16, 2026), https://www.csis.org/analysis/department-commerce-restricted-access-anthropics-latest-models-what-comes-next; Alex Zhao, *Is Access to Fable an Export?*, Harv. L. Rev. Blog (June 26, 2026), https://harvardlawreview.org/blog/2026/06/is-access-to-fable-an-export/.

b. The International Traffic in Arms Regulations (ITAR)

i. Model Weights as "Technical Data" and Related Defense Services

The ITAR, administered by the State Department's Directorate of Defense Trade Controls (DDTC), governs exports, reexports, retransfers, and temporary imports of defense articles on the United States Munitions List (USML), including technical data, as well as the furnishing of defense services to foreign persons.[40] "Technical data" includes information required for the design, development, production, operation, and other specified functions of a defense article.[41] Technical data may be recorded or stored in any form, so the fact that model weights exist as machine-readable files rather than human-readable documentation poses no definitional obstacle.[42] In addition, unlike the EAR, which classifies software and technology separately, the ITAR treats software directly related to defense articles as technical data.[43] Model weights could qualify as technical data if they, first, contain information required for a controlled defense function, or second, encode or reveal other ITAR-controlled technical data.[44] However, DDTC has not issued guidance addressing model weights as such, so both the first and second paths remain untested applications of existing definitions.

A "defense service" controlled by the ITAR can include assistance, such as training, provided to a foreign person in designing, producing, or using a defense article.[45] Separate from the status of model weights, a developer or downstream provider could furnish a defense service, for example, by using a model to assist a foreign person with an ITAR-controlled defense article.[46]

ii. The "Public Domain" Exclusion

The ITAR carves out certain publicly available information through its "public domain" provision, which is narrower than the EAR's "published" exclusion. To be eligible, information must be published and generally accessible through specified channels, including bookstores, unrestricted subscriptions, public libraries, patents, public conferences, approved public releases, and qualifying university research.[47] However, merely posting information on the internet is not included as an eligible channel.[48] Uploading ITAR-controlled technical data to the public internet therefore may or may not place it in the public domain and, if not, could be an unauthorized export. A curated repository might qualify as a public library,[49] but that theory remains legally

[40] 22 C.F.R. §§ 120.14(a)–(b), 120.31–.33, 120.57.
[41] 22 C.F.R. § 120.33(a).
[42] 22 C.F.R. § 120.31(a)(1).
[43] 22 C.F.R. § 120.33(a)(4).
[44] *See* 22 C.F.R. § 120.33(a). Model weights may also fall within the definition of a "defense article" in § 120.31, which includes items "that reveal technical data directly related to" USML items. Whether trained weights "reveal" the underlying ITAR-controlled technical data is fact-specific and turns on the extent to which the training data is recoverable from the weights.
[45] 22 C.F.R. § 120.32(a)(1).
[46] *See* 22 C.F.R. § 120.32(a)(1).
[47] 22 C.F.R. § 120.34(a).
[48] *Compare* 15 C.F.R. § 734.7(a)(4) (EAR: open list that includes "posting on the Internet on sites available to the public"), *with* 22 C.F.R. § 120.34(a) (ITAR: closed list with no internet posting provision).
[49] 22 C.F.R. § 120.34(a)(4).

DRAFT

uncertain. It may depend on whether the repository functions as a library—a collection assembled, organized, and made accessible by someone exercising judgment—rather than passively hosting what users choose to upload. In *Stagg v. U.S. Department of State*, the district court rejected the proposition that the internet as a whole is a public library while leaving open whether a genuine digital library may qualify.[50] DDTC, however, has never endorsed the idea that model repositories are public libraries and has historically read the public-domain provision narrowly.[51]

c. IEEPA and the ICTS Program as Alternative Authorities

The International Emergency Economic Powers Act (IEEPA) has not yet been used to broadly regulate transactions and services within the open-weight ecosystem. However, outside of export controls, IEEPA is the government's most realistic alternative authority for doing so.

Under IEEPA, the president can declare a national emergency based on an unusual and extraordinary threat that originates substantially outside the United States, and may then regulate transactions involving foreign property interests subject to U.S. jurisdiction.[52] A tailored executive order, for example, could restrict persons subject to U.S. jurisdiction from providing hosting, cloud, financial, or technical services to specified foreign model developers or platforms and could block property within U.S. reach. These measures could reach transactions and services involving foreign-origin models that fall outside the EAR, but without a U.S. nexus, they could not directly prevent a foreign developer acting abroad from publishing model weights. IEEPA also contains an informational-materials limitation, discussed below, that would pose a serious statutory obstacle to using the statute to prevent U.S. persons from distributing or downloading model weights.[53]

One existing IEEPA-based program relevant to downstream transactions in the open-weight ecosystem is the Information and Communications Technology and Services (ICTS) program.[54] Created under Executive Order 13873, the program is intended to protect the U.S. ICTS supply chain from technology and services developed or supplied by persons linked to a foreign adversary such as China.[55] Commerce may prohibit or mitigate transactions or classes of transactions involving such ICTS if those transactions have a U.S. jurisdictional nexus, involve a foreign interest, and present an undue or unacceptable risk. The covered technology categories include artificial intelligence, and an ICTS transaction can include the acquisition, importation, transfer, installation, dealing in, or use of ICTS, including the platforming or data hosting of

---

[50] *Stagg P.C. v. U.S. Dep't of State*, 354 F. Supp. 3d 448, 466 (S.D.N.Y. 2019), *vacated on other grounds*, 983 F.3d 589 (2d Cir. 2020) (finding that a "library open to the public" can include an internet library with digital collections but declining to construe the internet as a whole as a public library).

[51] *See* International Traffic in Arms: Revisions to Definitions of Defense Services, Technical Data, and Public Domain; Definition of Product of Fundamental Research; Electronic Transmission and Storage of Technical Data; and Related Definitions, 80 Fed. Reg. 31,525, 31,528 (June 3, 2015) (proposed rule) (proposing regulatory revisions to provide "a more explicit statement of the ITAR's requirement" that one must seek and receive a license or other authorization to release ITAR-controlled technical data prior to it becoming public domain).

[52] 50 U.S.C. §§ 1701(a), 1702(a)(1)(B).

[53] 50 U.S.C. § 1702(b)(3).

[54] *See* Exec. Order No. 13,873, §§ 1–2, 84 Fed. Reg. 22,689, 22,689–90 (May 15, 2019).

[55] *Id.* § 1(a), 84 Fed. Reg. at 22,689–90; 15 C.F.R. §§ 791.1–.4, 791.103, 791.109.

DRAFT

applications for consumer download.[56] Commerce therefore could restrict a U.S. cloud provider's hosting or distribution arrangement with a Chinese developer if the model weights or related service qualify as foreign-adversary ICTS and the other criteria are met.[57] The program's principal actions to date have focused on commercial actors rather than individual users and have emphasized the risk that a foreign-adversary supplier could retain access to or influence over the product, such as via software updates.[58] Part VI considers how the program could apply to U.S.-nexus distribution, hosting, and other downstream transactions involving model weights supplied by persons linked to a foreign adversary, and separately whether it could support direct restrictions on downloading or locally using stand-alone weight files.

d. Statutory and Constitutional Limits on Publication Controls

The "published" and "public domain" exclusions in the EAR and ITAR mitigate the potential for those regulatory regimes to conflict with the First Amendment. A license requirement imposed before a private party publishes technical information can operate as a prior restraint, and these exclusions reduce that concern by placing qualifying information outside the applicable licensing requirements. But the exclusions may not be broad enough to eliminate the concern, and the agencies could narrow them in the future. In fact, the ITAR's current definition of "public domain" was developed after the Justice Department concluded that an earlier version of the ITAR unconstitutionally restrained cryptographic disclosure.[59] In addition, the agencies' use of conduct-based controls, such as the U.S.-person controls under § 744.6, could implicate First Amendment concerns, regardless of whether the underlying information is published or in the public domain.

In other areas, the EAR exclusion for "published" material has some limits. Certain encryption software remains controlled even when publicly disclosed unless applicable conditions are met, such as required notifications; similarly, machine-ready firearms-production files remain controlled even when posted online.[60] The constitutionally required limits of those provisions, however, have not been firmly established. Litigation challenging export-control

---

[56] 15 C.F.R. § 791.4(a)(1) (designating the People's Republic of China, including Hong Kong and Macau, as a foreign adversary); *id.* § 791.3(a)(4)(v) (including artificial intelligence among the covered critical and emerging technologies); *id.* § 791.2 (defining "ICTS Transaction" to include the acquisition, importation, transfer, installation, dealing in, or use of ICTS, including "the platforming or data hosting of applications for consumer download," and providing that the term includes a class of ICTS transactions).

[57] 15 C.F.R. §§ 791.3(a), 791.103(a)–(b), 791.109(a)–(d).

[58] *See* Final Determination: Case No. ICTS-2021-002, Kaspersky Lab, Inc., 89 Fed. Reg. 52,434 (June 24, 2024) (prohibiting Kaspersky's provision and commercial distribution of cybersecurity and anti-virus software to persons in the United States); Securing the Information and Communications Technology and Services Supply Chain: Connected Vehicles, 90 Fed. Reg. 5,360 (Jan. 16, 2025) (prohibiting, by class rule, transactions by connected-vehicle manufacturers and hardware importers involving vehicle-connectivity hardware and covered software supplied by persons linked to China or Russia).

[59] In opinions issued between 1978 and 1984, the Department of Justice's Office of Legal Counsel concluded that the ITAR was unconstitutional insofar as it imposed a prior restraint on the disclosure of privately developed cryptographic information and recommended that the State Department narrow the regulations rather than await judicial correction. *See Bernstein v. U.S. Dep't of State*, 945 F. Supp. 1279, 1291–92 (N.D. Cal. 1996) (describing the opinions).

[60] *See* 15 C.F.R. §§ 734.7(b)–(c), 742.15(b); Control of Firearms, Guns, Ammunition and Related Articles the President Determines No Longer Warrant Control Under the United States Munitions List (USML), 85 Fed. Reg. 4,136 (Jan. 23, 2020).

DRAFT

restrictions on the online distribution of machine-ready firearm files has not produced a merits decision upholding those restrictions under the First Amendment.[61] Challenges to encryption rules resulted in decisions recognizing the expressive character of source code, but those cases involved code that programmers could read and use to communicate ideas.[62] Thus, these precedents do not fully clarify whether the First Amendment would permit continued control of model weights after their public release.

The ITAR narrows its "public domain" exclusion in a different way. As discussed above, § 120.34 does not treat ordinary internet posting as placing information in the public domain. A 2015 proposed rule would have expressly required authorization before public internet release, but DDTC did not finalize it.[63] In *Stagg*, a plaintiff brought constitutional challenges to the ITAR's licensing scheme, and the district court held that the ITAR did not require the plaintiff to obtain authorization to republish public domain information after aggregating it and making non-substantive modifications. The court rejected, however, the broader claim that any internet posting automatically confers public-domain status. The Second Circuit agreed with the district court's construction as applied to the conduct alleged in the complaint and directed dismissal because that construction of the regulations mooted the constitutional claims.[64] As with the EAR, the case law thus leaves unresolved whether the First Amendment would permit the ITAR to require authorization before the online release of model weights.

Before any constitutional question is reached, an IEEPA or ECRA item-based control would confront a separate statutory limit. The Berman Amendment is a provision within IEEPA that withholds authority to regulate, directly or indirectly, the import or export of "information or informational materials," and ECRA carries that limitation into its item-based controls, subject to an exception for certain exports controlled for nonproliferation or antiterrorism purposes.[65] As applied to model weights, two questions arise. The first is whether weights are "information or informational materials" under the statute. The non-exclusive examples cited in IEEPA—

---

[61] *See Defense Distributed v. U.S. Dep't of State*, 838 F.3d 451, 458–61 (5th Cir. 2016) (affirming denial of a preliminary injunction on balance-of-equities and public-interest grounds without deciding the First Amendment merits), *cert. denied*, 138 S. Ct. 638 (2018); *Washington v. U.S. Dep't of State*, 996 F.3d 552, 561–64 (9th Cir. 2021) (holding that ECRA and the AECA precluded the plaintiff States' administrative challenges).

[62] *See Junger v. Daley*, 209 F.3d 481, 484–85 (6th Cir. 2000) (holding that encryption source code is an expressive means for exchanging information and ideas about computer programming); *Bernstein v. U.S. Dep't of State*, 922 F. Supp. 1426, 1434–36 (N.D. Cal. 1996) (concluding that cryptographic source code is a language used by cryptographers to communicate); *see also Universal City Studios, Inc. v. Corley*, 273 F.3d 429, 445–51 (2d Cir. 2001) (recognizing that computer code may combine expressive and functional elements).

[63] *See* International Traffic in Arms: Revisions to Definitions of Defense Services, Technical Data, and Public Domain; Definition of Product of Fundamental Research; Electronic Transmission and Storage of Technical Data; and Related Definitions, 80 Fed. Reg. 31,525, 31,528 (June 3, 2015) (proposed rule).

[64] *Stagg P.C. v. U.S. Dep't of State*, 354 F. Supp. 3d 448, 461–66 (S.D.N.Y. 2019), *vacated with instructions to dismiss*, 983 F.3d 589, 602–09, 613 (2d Cir. 2020).

[65] 50 U.S.C. § 1702(b)(3); *id.* § 4813(b) (providing that ECRA authority may not be used to regulate or prohibit the export, reexport, or in-country transfer of any item that § 1702(b)(3) places beyond IEEPA). ECRA exempts functions exercised under the statute from the judicial-review provisions of the Administrative Procedure Act (APA), except for certain enforcement-related actions. 50 U.S.C. § 4821(a). Because ordinary APA review is unavailable, a pre-enforcement challenge to an agency's statutory authority would need to proceed as an *ultra vires* claim, which would succeed only if the plaintiff could show that the agency plainly acted in excess of its delegated powers and contrary to a clear and mandatory statutory requirement. *See Fed. Express Corp. v. U.S. Dep't of Com.*, 39 F.4th 756, 763–66 (D.C. Cir. 2022). Section 4821(a) does not, however, bar constitutional claims.

DRAFT

including publications, films, posters, photographs, tapes, and news wire feeds—generally concern expressive or communicative works or media, while weights are numerical parameters transferred largely for their functional capabilities. Recent Justice Department interpretations distinguish protected expressive material from technical or functional data.[66] On the other hand, the statute is broad and applies "regardless of format or medium of transmission," and even if only expressive materials are protected, completed weights may still qualify—a related argument taken up in the discussion of the First Amendment analysis below.

Second, even if model weights qualify as "information or informational materials," the exemption has been understood to protect the exchange of informational materials that already exist, but not the services that bring the materials into being or alter them.[67] An argument could therefore be made that the Berman Amendment covers completed model weights but services like model training or customization fall outside its scope, at least based on existing regulatory constructions of the exemption. The Berman Amendment may also reach service restrictions that indirectly prohibit the exchange of protected materials. In the 2020 *TikTok* litigation, courts preliminarily enjoined IEEPA-based prohibitions on hosting and related platform services that would have prevented users from exchanging videos and other informational materials across borders, which the courts determined likely constituted indirect regulation forbidden by the Berman Amendment.[68]

While the statutory and constitutional inquiries cannot be collapsed into one analysis—indeed, the Justice Department expressly rejected wholesale incorporation of First Amendment doctrine into the Berman Amendment[69]—both inquiries may nevertheless depend in part on whether model weights are expressive or predominantly functional. The First Amendment analysis would begin by asking whether the weights are protected expression, or, even if not, whether the restriction disproportionately burdens other protected First Amendment activities. On the one hand, model weights are numerical parameters that enable software to produce outputs. Some scholarship argues that weights are predominantly functional rather than expressive.[70] Unlike source code, which programmers can read as a symbolic language, model weights are generally not intelligible through direct review and are used primarily by software

[66] *See* 28 C.F.R. § 202.226(a) (construing, for purposes of the IEEPA-based Data Security Program, the Berman Amendment to protect "expressive material" but not "technical, functional, or otherwise non-expressive data"). The executive branch's construction of Berman would inform a court's analysis, but the reviewing court would independently determine its meaning. *See Loper Bright Enters. v. Raimondo*, 603 U.S. 369, 394–95 (2024).

[67] *See* 28 C.F.R. § 202.226(b)(1) (excluding material not fully created at the time of the transaction, substantive or artistic alteration or enhancement, and marketing and business consulting services used to create such material); 31 C.F.R. § 560.210(c)(2) (similar); *see also United States v. Amirnazmi*, 645 F.3d 564, 585–88 (3d Cir. 2011).

[68] *See TikTok Inc. v. Trump*, 507 F. Supp. 3d 92, 108–09 (D.D.C. 2020) (concluding that prohibitions on internet-hosting and related platform services would indirectly restrict the cross-border exchange of videos, photographs, and other informational materials); *Marland v. Trump*, 498 F. Supp. 3d 624, 636–41 (E.D. Pa. 2020) (similar). Both decisions arose at the preliminary injunction phase and rested at least in part on conclusions that the plaintiffs were likely to succeed under IEEPA's statutory limitation—but neither was a final merits determination.

[69] *See* Preventing Access to U.S. Sensitive Personal Data and Government-Related Data by Countries of Concern or Covered Persons, 90 Fed. Reg. 1,636, 1,674 (Jan. 8, 2025).

[70] *See* Alan Z. Rozenshtein, *There Is No General First Amendment Right to Distribute Machine-Learning Model Weights*, Lawfare (Apr. 4, 2024, 8:02 AM), https://www.lawfaremedia.org/article/there-is-no-general-first-amendment-right-to-distribute-machine-learning-model-weights; Julia Trevino & Anita Nikolich, *Are AI Model Weights Protected Speech Under the First Amendment?* (Mar. 26, 2026) (unpublished manuscript), https://ssrn.com/abstract=6538918.

DRAFT

during inference. Although encryption-era case law establishes that the First Amendment protects some code that communicates ideas, especially in source-code form, not every machine-readable file is necessarily speech.[71] The Third Circuit's 2026 decision in *Defense Distributed* likewise required a fact-specific inquiry into whether a particular use of code burdened by a regulation involved the expression or communication of ideas.[72]

On the other hand, humans exercise judgment when they create an AI model, and the resulting model weights often reflect the developers' creative choices. Some developers now shape a model's character by authoring a written "constitution," a document describing the values the model should hold, which guides the model weights during the training process.[73] A model trained this way converts its developer's moral values into a functional medium, and what the model will say or refuse in its outputs can be traced in part to a text that is clearly expressive. In that respect, the weights resemble source code, which courts have recognized can embody expression despite its functional character. But the fact that weights reflect human choices does not automatically convert them into protected expression; what matters is whether the particular use of the weights that is being burdened by the regulation involves the expression or communication of ideas.[74] Because training methods differ, the strength of this argument will vary by model; weights trained to embody a published statement of values present a stronger expressive claim than weights shaped only by efforts to optimize how the model performs on benchmarks. Some scholars also argue that AI-generated outputs may be protected through the rights of model users,[75] supporting a related claim that restricting the weights would burden users' access to protected information. Whether or not the weights themselves are found to be expressive, preventing a developer from publishing the weights could disproportionately burden the protected interests of users.[76]

If the First Amendment applies, courts would assess whether any pre-publication licensing regime includes clear standards that would limit agency discretion and whether any restriction is tailored to the risk.[77] National-security interests would carry substantial weight, but they would not eliminate the need to implement procedural safeguards. With respect to tailoring, a restriction would be difficult to justify when the weights of materially comparable foreign

---

[71] *See Universal City Studios, Inc. v. Corley*, 273 F.3d 429, 445–51, 454–55 (2d Cir. 2001).
[72] *Defense Distributed v. Att'y Gen. of N.J.*, 167 F.4th 65, 93–99 (3d Cir. 2026).
[73] Anthropic, *Claude's Constitution* (Jan. 2026), https://www.anthropic.com/constitution; Anthropic, *Claude's New Constitution* (Jan. 22, 2026), https://www.anthropic.com/news/claude-new-constitution; *see also* Doğa Özden, *The First Amendment Argument Anthropic Didn't Make*, Dorf on Law (Mar. 16, 2026), https://www.dorfonlaw.org/2026/03/the-first-amendment-argument-anthropic.html.
[74] *See Defense Distributed v. Att'y Gen. of N.J.*, 167 F.4th 65, 93–99 (3d Cir. 2026) (performing a "fact-based and conduct-specific analysis" of the code, focusing on the technical nature of the code, how the code is used in context, who is communicating through the code and the intended recipient of the communication, for what purpose or purposes the code operates, and what, if anything, the code communicates).
[75] *See, e.g.,* Eugene Volokh, Mark A. Lemley & Peter Henderson, *Freedom of Speech and AI Output*, 3 J. Free Speech L. 651, 651–59 (2023).
[76] *See Arcara v. Cloud Books, Inc.*, 478 U.S. 697, 703–07 (1986) (recognizing that First Amendment scrutiny has applied to some statutes directed at non-expressive activity that disproportionately burdens protected First Amendment activities); *see also TikTok Inc. v. Garland*, 604 U.S. 56, 67–69 (2025) (assuming without deciding that restrictions on foreign-adversary control of a communications platform fell within that category and were subject to First Amendment scrutiny).
[77] *See City of Lakewood v. Plain Dealer Publ'g Co.*, 486 U.S. 750, 757–58 (1988); *Freedman v. Maryland*, 380 U.S. 51, 58–59 (1965).

models have already been published and are readily available, as the burden on speech would not meaningfully reduce foreign adversaries' access to the relevant capabilities.[78] A post-release rule would be even harder to justify because copies of the model would likely already be beyond recall. The government's strongest case would likely involve a narrow pre-release measure for a model posing a serious national security risk with no comparable foreign substitute. A narrowly drawn restriction may therefore be possible, even if the First Amendment applies, but its validity would depend heavily on the facts and the design of the rule.

e. AI Model Outputs as Controlled Exports

Beyond the export-control status of model weights, AI-generated outputs require a separate analysis. An output may constitute EAR technology or software or ITAR technical data when its content meets the relevant definition. Those definitions do not depend on whether the information was produced by a person or a model.[79] A source-code file generated by a model may qualify as controlled software, and a technical explanation provided by a model may qualify as controlled technology or technical data. The same is true whether the output comes from a closed-weight or open-weight system.

Publishing a model's weights and generating outputs from that model are distinct events that require separate assessments of whether they constitute a release of controlled technology, software, or technical data. Publishing a model's weights does not confer "published" or "public domain" status on outputs produced by that model; each output must qualify for the EAR's § 734.7 exclusion or the ITAR's § 120.34 provision on its own terms. The inquiry should instead focus on the information disclosed and the transaction through which it is disclosed. A model's restatement of information already published or in the public domain ordinarily retains that status. Conversely, information does not qualify as published merely because a model produced it for a user, as the act of generating such an output, without more, does not result in unrestricted public dissemination.[80] A model trained on public sources may also synthesize or infer information not found in any single source.[81] Such a synthesis can be controlled if its content amounts to controlled technology, software, or technical data and its disclosure is a transaction the regulations reach, such as a release to a foreign person.

Both closed and open-weight models can generate controlled outputs, but they differ in how readily safeguards can be applied to reduce that risk. For a closed model, the deployer can apply safeguards such as account controls, logging, and monitoring to reduce the disclosure of controlled outputs to foreign-person users. Because the developer of an open-weight model cannot control the deployment context—including whether these safeguards are applied—the allocation of export-compliance responsibility for open model outputs becomes especially important.

DRAFT

---

[78] *See infra* note 118 (discussing the model-weights foreign direct product rule and its implications).
[79] LRI Report, *supra* note 13, at 8; *see also* Joe Khawam & Tim Schnabel, *AI Model Outputs Demand the Attention of Export Control Agencies*, Just Security (Dec. 12, 2025), https://www.justsecurity.org/126643/ai-model-outputs-export-control/ [hereinafter Khawam & Schnabel, *AI Model Outputs*].
[80] LRI Report, *supra* note 13, at 11.
[81] *Id*.

IV. Allocation of Export Control Responsibilities Across the Open-Weight Lifecycle

Part III of this paper assessed how model weights and outputs, along with related conduct, are controlled under the EAR or ITAR. This Part now asks which actor is primarily responsible for compliance at different stages of the open-weight lifecycle. Under the EAR, the exporter is generally the person in the United States with "the authority of a principal party in interest to determine and control the sending of items out of the United States."[82] The ITAR does not use the same defined term, but it also places authorization responsibility on the person who controls the export.[83] The responsibility for releasing model weights and disclosing model outputs can therefore be assigned to different actors as control over each event shifts. Item-based obligations, including the item-based catch-all controls, fall away once the weights qualify as published, while conduct-based obligations, such as the U.S.-person controls, attach to the actor performing the conduct and apply regardless of the weights' publication status.[84]

a. Developer Responsibilities

The original developer of a model most clearly controls the initial release of the weights. Under the EAR, the developer must determine whether the weights will qualify as "published" under § 734.7. Weights that are unclassified (meaning that they are not classified national security information) and are made public without restrictions on further dissemination generally fall outside the EAR.[85] A permissive open-source license most clearly meets this standard, while a model-specific license with more substantive restrictions might not.

Weights that do not qualify as "published" remain subject to the EAR if they otherwise fall within its jurisdiction. If no CCL entry applies, the weights will generally be designated EAR99, which leaves them subject to the EAR but without a list-based license requirement. For such weights, an export may still require a license under other controls (e.g., destination-based, end-use, end-user, or embargo), and BIS may also inform a specific party under an item-based catch-all provision, such as § 744.22(b), that a particular transaction requires a license.[86] Moreover, despite Commerce's announced policy of not enforcing the AI Diffusion Rule, some of the text remains in the EAR, including ECCN 4E091.[87] For a developer whose model approaches the 10^26 threshold, the entry's continued presence complicates compliance. Pre-release transfers of unpublished frontier weights fall within regulatory text that on its face requires a license, and an announced enforcement policy, unlike a rescission, can be reversed without further rulemaking. A cautious developer may therefore treat the entry as operative until formally removed.

The ITAR poses a narrower but more challenging issue for developers. Weights developed for a defense article on the U.S. Munitions List, or for a specified function of such an article, may constitute technical data, and the same may be true of weights that were trained on

[82] 15 C.F.R. § 772.1 (defining "exporter").
[83] *See* 22 C.F.R. §§ 120.15, 120.50, 122.1, 123.1, 123.9, 123.22.
[84] *See* 15 C.F.R. § 744.6.
[85] 15 C.F.R. § 734.7(a)(4).
[86] *See, e.g.,* 15 C.F.R. pts. 744, 746; *id.* § 744.22(b).
[87] *See supra* Part III.a.iii.

DRAFT

controlled technical data, if the data can be recovered from the weights.[88] Because the ITAR does not consider ordinary internet posting to qualify as a public-domain channel, a developer cannot rely on posting alone as a compliance strategy.[89] But in practice, the weights of such models are unlikely to be released, since contractual restrictions and a desire to protect trade secrets would often prevent the developer from releasing the weights regardless of export controls.

Finally, a developer weighing an open release should recognize that its obligations do not end at publication. Section 744.6 can restrict U.S. persons from providing certain types of support whether or not the weights remain subject to the EAR.[90] Criminal prohibitions on knowingly assisting biological- or chemical-weapons activities would apply the same way, as would sanctions on dealings with specified parties.[91] These obligations attach to the developer's conduct rather than to the weights, so they persist after control over the weights is lost.

b. Downstream Releases, Deployments, and Infrastructure

Downstream actors are responsible for new transfers that they control. A person that fine-tunes published weights and redistributes the derivative model must analyze that release as a separate event rather than assuming the model inherits the status of the original release. A fine-tuned version of a previously published model that itself is posted on the internet without restrictions on its dissemination may still benefit from the EAR exclusion. On the other hand, redistributing it in a selective manner or with restrictive access terms could produce a different result. The ITAR could be implicated if the derivative model has a defense-specific design or is fine-tuned on ITAR-controlled technical data. Fine-tuners and other downstream actors are also subject to the same conduct-based obligations discussed above.[92]

A hosted inference provider deploys an open-weight model to offer users access without transferring the weights. BIS has treated the provision of cloud capacity and the remote use of hosted software, without download, as services rather than exports.[93] However, the June 2026 letter to Anthropic adds uncertainty because it appears to conflate access to a hosted model with an export, such that licenses would be required for foreign access to the models. BIS did not explain how it would apply this rationale as a general rule or distinguish its prior guidance on cloud services, so providers face an uncertain enforcement risk until BIS clarifies its position.[94] Regardless, the EAR may still apply to some aspects of providing hosted inference, such as where a provider discloses controlled model outputs or provides prohibited assistance.

The same considerations apply to a commercial vendor that integrates an open-weight model into a hosted service that does not involve transferring the weights. If the service delivers model-generated technical information to foreign users, the vendor controls those disclosures and

DRAFT

[88] 22 C.F.R. §§ 120.31, 120.33(a), 121.1.
[89] 22 C.F.R. § 120.34(a).
[90] 15 C.F.R. § 744.6.
[91] 18 U.S.C. §§ 175, 229; 31 C.F.R. ch. V.
[92] *See* 15 C.F.R. § 744.6; 31 C.F.R. ch. V; 18 U.S.C. §§ 175, 229.
[93] *See* sources cited *supra* note 38.
[94] *See supra* Part III.a.v.

is responsible for export-control compliance. By contrast, if model weights are integrated into a device, shipping that device would also transfer (and potentially export) the weights.

Infrastructure-as-a-service (IaaS) providers are further removed from responsibility for exports. Unlike a hosted inference provider, which deploys a model and offers it to users, an IaaS provider supplies computing capacity on which customers deploy and run models of their choosing. IaaS providers generally supply compute and are not exporters of data that customers create or transmit.[95] They may nonetheless face secondary liability if they know or have reason to know that customers are using their services to violate the EAR.[96] BIS has proposed customer-identification requirements and issued red-flag guidance relevant to advanced-computing transactions and AI model training.[97] An IaaS provider's compliance role therefore centers on conducting due diligence, screening, maintaining records, and responding to known risks rather than on classifying every model output generated on customer infrastructure.

Universities and other research institutions face a distinct set of questions. The EAR generally excludes technology and software that result from fundamental research intended for publication.[98] This exclusion applies only to the *results* of qualifying research, not the inputs used to produce them. Thus, when a university develops an open-weight model, controlled training data used in the project remains controlled, and providing that data to foreign-person researchers during the work can require authorization from BIS even when the eventual results of the research will be published. Whether a trained model qualifies as a research result may depend on the project and publication plan.

A foreign person who downloads weights that are outside the EAR because they were published does not thereby acquire EAR obligations, and reexport controls do not attach to the weights. But publishing one item does not remove controls that may apply to the rest of the deployment stack (e.g., U.S.-origin compute, controlled training data or software) or to later transactions with a U.S. nexus.

c. Responsibility for AI-Generated Outputs

In assessing how export controls apply to model outputs, the relevant event is the disclosure of controlled technology or technical data to a foreign person. When a system operator with a U.S. nexus permits a foreign person to access the outputs of a model, that operator is the

---

[95] *See* sources cited *supra* note 38.
[96] *See* 15 C.F.R. § 764.2(b).
[97] *See* Taking Additional Steps To Address the National Emergency With Respect to Significant Malicious Cyber-Enabled Activities, 89 Fed. Reg. 5,698 (Jan. 29, 2024) (proposed rule) (to be codified at 15 C.F.R. pt. 7); Bureau of Indus. & Sec., U.S. Dep't of Com., *BIS Policy Statement on Controls That May Apply to Advanced Computing Integrated Circuits and Other Commodities Used to Train AI Models* (May 13, 2025), https://www.bis.gov/media/documents/ai-policy-statement-training-ai-models-may-13-2025; 15 C.F.R. pt. 732, supp. no. 3 (BIS's "Know Your Customer" Guidance and Red Flags). General Prohibition 10 provides a related example of this knowledge-based structure at the compute layer. BIS has advised that the use of certain PRC advanced-computing ICs, including Huawei Ascend chips, may implicate General Prohibition 10 where a person has knowledge that an EAR violation has occurred, is about to occur, or is intended in connection with an item subject to the EAR. *See* 15 C.F.R. § 736.2(b)(10); Bureau of Indus. & Sec., U.S. Dep't of Com., *Guidance on Application of General Prohibition 10 (GP10) to People's Republic of China (PRC) Advanced-Computing Integrated Circuits (ICs)* 1–2 (May 13, 2025), https://www.bis.gov/media/documents/general-prohibition-10-guidance-may-13-2025.pdf.
[98] 15 C.F.R. § 734.8(a)–(c).

DRAFT

most plausible exporter because it controls access.[99] However, no agency has explicitly addressed how responsibility should be allocated for AI-generated outputs, so this analysis remains the most defensible reading of existing definitions rather than settled doctrine.

A developer that only publishes weights and does not control outputs from an independent downstream deployment ordinarily should not be treated as the exporter of later outputs from that model. The developer's relevant act was only releasing the weights. It also remains subject to its own conduct-based obligations (such as the § 744.6 prohibitions on U.S.-person support), but it is not the exporter of outputs generated when a third party independently uses the model. By contrast, a developer that also provides access to the model via a public demo or API is likely responsible for the outputs released through that service.[100] Likewise, a third-party host or commercial integrator is responsible for outputs it releases through its service. Those actors control a number of levers for complying with export controls, such as screening their customers, logging use of the system, and monitoring abuse. Moreover, any U.S. entity may face deemed-export risks if foreign-national employees or contractors receive export-controlled outputs from an internally deployed model. Technology Control Plans may be appropriate in such circumstances.[101]

Each of the scenarios above involves an actor with a U.S. nexus. The framework, however, reaches its limits when there is none. If a foreign person downloads published weights, runs them on foreign hardware, and generates information abroad, no one exports that model output from the United States.[102] To conclude otherwise would be to treat the original publisher as the exporter of every later output the model produces—a theory that would effectively nullify the "published" exclusion by making publication a source of perpetual responsibility for third parties' downstream use.[103] In that wholly foreign scenario, foreign law may regulate the conduct, but U.S. export controls do not. By contrast, a U.S. person who runs an open-weight model locally and sends controlled outputs to a foreign person effects an export.

V. Policy Considerations for Open-Weight Models

DRAFT

[99] LRI Report, *supra* note 13, at 8–9.
[100] *Id.* at 9 n.22.
[101] *Id.* at 13–16; Joe Khawam & Tim Schnabel, *Vetting Foreign AI Talent: Security Without Exclusion*, Just Security (July 8, 2026), https://www.justsecurity.org/145569/vetting-foreign-ai-talent-security-without-exclusion/. The same internal-use challenge can arise for open-weight developers, redistributors, fine-tuners, and hosted providers that give foreign-person employees or contractors access to internal deployments, pre-release or modified models, evaluation systems, prompts, outputs, or remediation discussions capable of revealing controlled technology or technical data. *See* Khawam & Schnabel, *AI Model Outputs*, *supra* note 79.
[102] *See* 15 C.F.R. § 734.13(a)(1)–(2) (defining "export" to include an actual shipment or transmission out of the United States and the release of "technology" or source code to a foreign person in the United States); 22 C.F.R. § 120.50(a) (defining "export" under the ITAR to include, among other things, an actual shipment or transmission out of the United States and the release of technical data to a foreign person in the United States).
[103] Foreign direct product rules do not provide an alternative basis for jurisdiction over outputs generated abroad. Each rule applies only when its specified technology, software, product, destination, and end-user conditions are met, and outputs derived from published weights would not satisfy them. The weights, once published, are not subject to the EAR and are not the specified technology or software of any entry. *See* 15 C.F.R. § 734.9.

By keeping a model closed, a developer can directly monetize access to it and keep trade secrets confidential. Additionally, because closed-weight models remain under the developer's control, the developer can monitor abuse and deploy (and update) safeguards. By contrast, a developer that releases model weights enables users to deploy the model locally, customize it, and conduct research on it directly, but the developer cannot directly monetize it as easily.[104] Firms may also combine the two approaches by keeping the most capable system closed while releasing smaller or earlier versions. The United States has an interest in remaining competitive in both modes, particularly because ceding the open-weight market to China would not stop new capabilities from being available.

The National Telecommunications and Information Administration's (NTIA) July 2024 report remains the U.S. government's most thorough analysis of the risks and benefits of open-weight models.[105] NTIA analyzed the marginal risk that such models pose beyond the alternatives that are already available (e.g., closed APIs or other open models, as well as non-AI tools).[106] Applying this lens to export controls means asking whether restricting the release of model weights meaningfully denies or delays an adversary's access to a valuable capability, and even if so, whether the risks posed by that access would outweigh the benefits of publicly available weights. The report recommended against restricting open-weight releases "at this time," but also recommended bolstering the government's ability to monitor whether future models pose risks that merit limiting model weight availability.

NTIA identified important benefits from open-weight models. They can reduce costs, thereby lowering barriers to entry in some fields. They can fuel competition in developing specialized systems. They can enable research, including independent testing and reproducibility of AI research. They also permit local deployment for privacy or security reasons.[107] Wide adoption of U.S.-origin open models can also support efforts to set global AI norms, and allies and partners with access to them may rely less on models from countries of concern.[108]

The report also identified risks. As discussed above, open-weight models can be modified, stripped of safeguards, and deployed in contexts that prevent monitoring. They may lower barriers to using AI for cyberattacks, CBRN programs, fraud, and military or intelligence uses. Because published weights can no longer be controlled, an error in assessing whether an open-weight model would advance the frontier of capabilities may have lasting consequences.

---

[104] OpenAI, *OpenAI's Comment to the NTIA on Open Model Weights* (Mar. 27, 2024), https://openai.com/global-affairs/openai-s-comment-to-the-ntia-on-open-model-weights/; Caroline Meinhardt et al., *Beyond DeepSeek: China's Diverse Open-Weight AI Ecosystem and Its Policy Implications*, Stanford Inst. for Human-Centered AI (Dec. 16, 2025), https://hai.stanford.edu/policy/beyond-deepseek-chinas-diverse-open-weight-ai-ecosystem-and-its-policy-implications.

[105] NTIA Report, *supra* note 9.

[106] *See id.* at 10–11.

[107] A concrete example of the value of local open-weight model deployment occurred in July 2026, when an autonomous agent powered by OpenAI models escaped a cyber-evaluation sandbox and compromised Hugging Face's infrastructure. When safety guardrails on commercial frontier closed models blocked Hugging Face's analysis of the attack logs—given the models' inability to distinguish an incident responder from an attacker—Hugging Face conducted the forensic review on its own infrastructure using GLM 5.2, a Chinese open-weight model. *See* Hugging Face, *Security Incident Disclosure—July 2026* (July 16, 2026), https://huggingface.co/blog/security-incident-july-2026.

[108] *See* NTIA Report, *supra* note 9, at 22–23.

DRAFT

Nevertheless, NTIA found the evidence insufficient to definitively determine either that restrictions on such open-weight models are warranted, or that restrictions will never be appropriate in the future.[109]

The approach taken in the AI Diffusion Rule reflected the marginal-risk analysis of the NTIA report. It applied export controls to some unpublished model weights while excluding published model weights and closed models less capable than the strongest published model.[110] Those exclusions were consistent with the marginal risk that those models posed. If such a model's weights were stolen or transferred without authorization, the recipient would gain nothing beyond the capabilities that the strongest open model already provides to anyone, so a license requirement would impose costs without denying additional capabilities. BIS acknowledged that unilateral controls on behind-frontier models would be ineffective when equivalent or superior systems are available abroad.[111] The rule thus treated unpublished frontier weights as a more administrable chokepoint than open publication of weights. It also showed that BIS viewed foreign availability as highly relevant to whether controls should be imposed on models subject to U.S. jurisdiction.

Assessing the foreign availability of open-weight models should be central to any open-weight policy. Chinese open-weight releases, including in the DeepSeek, Qwen, Kimi, GLM, and MiniMax families, have become widely used and are increasingly capable. The Chinese open-weight ecosystem is not only diverse, including general-purpose and specialized models from multiple firms, but is highly consequential given substantial adoption of those models.[112] As CSIS has argued, semiconductor controls can slow Chinese progress but cannot halt it, especially to the extent that algorithmic and systems innovation can offset hardware constraints.[113] But the mere existence of Chinese open models does not justify relaxing controls on U.S. model weights unless the Chinese models actually supply the capabilities that U.S. export control would be designed to deny.

Because general-purpose models support many tasks, foreign availability should be assessed capability by capability rather than model by model—i.e., whether a Chinese model matches U.S. systems on the specific functions a control is meant to deny, such as cyber operations or biological assistance. Overall benchmark scores are thus less useful, as they aggregate AI models' performance across many tasks and do not reveal gaps in the critical areas. Moreover, foreign availability assessments need to take into account how models can be deployed. A foreign model may not serve as an adequate substitute for a restricted U.S. model if the relevant actors cannot run the foreign model on the hardware available to them. When a Chinese open model offers comparable capability in relevant areas at similar cost, restricting the release of U.S.-origin weights would do little to deny adversaries the capability at issue but

DRAFT

[109] *See id.* at 3, 47.
[110] *See* AI Diffusion Rule, *supra* note 21, at 4,553–56, 4,578–79.
[111] *Id.* at 4,553–55.
[112] Meinhardt et al., *supra* note 104.
[113] Gregory C. Allen, *DeepSeek, Huawei, Export Controls, and the Future of the U.S.-China AI Race*, Ctr. for Strategic & Int'l Stud. (Mar. 7, 2025), https://www.csis.org/analysis/deepseek-huawei-export-controls-and-future-us-china-ai-race.

would instead tend to push users toward the Chinese ecosystem. Influence over safety and deployment practices would also migrate with them.

Yet the capabilities of open-weight models are close enough to those of closed-weight frontier models that the absolute risk associated with open-weight releases will continue to grow. Epoch AI's capability index estimates that leading open models have trailed frontier models by about four months since the start of 2026.[114] The aggregate figure, however, misses the variation that likely exists across domains, and for assessing risk, the relevant measure is the gap in security-relevant capabilities. The U.K. AI Security Institute found that open-weight models lag four to seven months behind on cyber evaluations.[115] The gap between closed and open-weight models has narrowed in part due to distillation, in which a weaker student model is trained on the outputs of a stronger teacher model.[116] As the frontier advances, open systems can follow close behind.

Those comparisons, however, measure the open ecosystem against the overall frontier; they do not establish the marginal contribution of any particular U.S. release, which depends on the recipient's realistic alternatives—particularly comparable foreign open models. Recent public evaluations suggest that Chinese models constitute much of the current open-weight frontier in domains relevant to export controls, while at least one major U.S. open-weight release did not substantially advance that frontier.[117] Foreign developers are likely to continue releasing ever-more-capable open models that, once their weights are published, can be deployed largely outside the reach of U.S. export controls.[118] The current market therefore combines rising absolute risk from open models with evidence that U.S. releases may add little risk at the margin.

---

[114] Jack Edwards & Luke Emberson, *Open Models Lag State-of-the-Art Closed Models by 4 Months*, Epoch AI (2026), https://epoch.ai/data-insights/open-closed-eci-gap.

[115] AI Sec. Inst., *How Far Behind the Frontier Are Leading Open Weight Models on Cyber?* (July 17, 2026), https://www.aisi.gov.uk/blog/how-far-behind-the-frontier-are-leading-open-weight-models-on-cyber.

[116] *See* Americans for Responsible Innovation, *Explainer: DeepSeek, Distillation, and AI IP Theft* (Mar. 30, 2026), https://ari.us/policy-bytes/explainer-deepseek-distillation-and-ai-ip-theft/ (explaining the mechanics of distillation); Dwarkesh Patel, *The Sample-Efficiency Black Hole* (2026), https://www.dwarkesh.com/p/the-sample-efficiency-black-hole ("I think the reason it is relatively easy for open source and previous laggards to catch up to within months of the frontier is that data is the real driver of progress. And data can be easily distilled from public APIs, whereas hyper-parameters and training tricks and architectural micro-optimizations cannot – if the latter were driving most of progress, then catching up would be harder than we are observing it to be.").

[117] *See* AI Sec. Inst., *How Far Behind the Frontier Are Leading Open Weight Models on Cyber?* (July 17, 2026), https://www.aisi.gov.uk/blog/how-far-behind-the-frontier-are-leading-open-weight-models-on-cyber (reviewing the cyber capabilities of GLM-5.2 and DeepSeek V4-Pro); Eric Wallace et al., *Estimating Worst-Case Frontier Risks of Open-Weight LLMs* 1–5 (Aug. 13, 2025), https://arxiv.org/abs/2508.03153 (comparing gpt-oss with DeepSeek R1-0528, Kimi K2, and Qwen3 Thinking).

[118] The AI Diffusion Rule created a model-weights foreign direct product rule (FDPR) that brings certain foreign-produced weights specified in ECCN 4E091 within the EAR when they are produced using certain advanced computing equipment subject to the EAR. 15 C.F.R. § 734.9(l). However, it would have a limited impact on the weights available from Chinese open model developers because Note 1 to 4E091 excludes parameters that have been "published" as defined in § 734.7(a), and BIS has explained that the rule did not impose controls on open-weight models. *See* ECCN 4E091 note 1, 15 C.F.R. pt. 774, supp. no. 1; *Framework for Artificial Intelligence Diffusion*, 90 Fed. Reg. 4,544, 4,554–55 (Jan. 15, 2025). Moreover, even where the unpublished weights of a foreign developer are captured by the rule, its practical effect on such a developer that disregards U.S. requirements—as many Chinese developers likely would—depends on whether BIS could reach an intermediary (e.g., a cloud provider, chip supplier, or some other commercial partner) or some other chokepoint.

DRAFT

Despite the rising overall risk, U.S. policy has been consistent in not restricting the publication of open weights. The second Trump administration decided not to enforce the AI Diffusion Rule,[119] but has not replaced it with an alternative rule. In July 2025, the administration issued an AI Action Plan that encourages open-source and open-weight AI, and the plan calls for stronger semiconductor enforcement, allied coordination, frontier-risk evaluation, and investments in biosecurity.[120] President Trump also signed Executive Order 14320, promoting export of the American AI technology stack.[121]

At the same time, however, the government began building selective oversight at the closed frontier. Executive Order 14409 directs agencies to design a voluntary pre-release review process for covered frontier models and expressly disclaims mandatory licensing of the models.[122] The administration also intervened in closed-model deployment, both by issuing the export-control directive to Anthropic covering its Mythos 5 and Fable 5 models, which led the company to restrict foreign-person access, and by asking OpenAI to limit the initial availability of GPT-5.6.[123] To date, these interventions have stopped at the closed frontier; none has reached the publication of model weights.

VI. Analysis and Recommendations

Parts III and IV of this paper identified the existing export-control authorities that apply to open-weight models and analyzed which actors would be held responsible for exports under those rules. Part V then showed that the practical value of those legal authorities depends on whether export controls actually reduce the access of adversaries to the relevant capabilities in light of the alternatives. This Part now applies that analysis separately to model weights and then to model outputs and downstream conduct to identify where export controls can materially reduce risk in the open-weight ecosystem and where they cannot.

a. Applying the Marginal-Risk Framework to the Release of Model Weights

Since NTIA's 2024 recommendation against restricting open-weight releases,[124] open models have become more capable and closer to the frontier, increasing the absolute risk associated with release. But the marginal risk of U.S. open-weight models depends on whether those releases would provide adversaries with access to capabilities in critical domains beyond

---

[119] Press Release, Bureau of Indus. & Sec., U.S. Dep't of Com., *Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls* (May 13, 2025), https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens.

[120] Exec. Off. of the President, Off. of Sci. & Tech. Pol'y, *Winning the Race: America's AI Action Plan* 7–8, 21–22 (July 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf [hereinafter AI Action Plan].

[121] Exec. Order No. 14,320, Promoting the Export of the American AI Technology Stack, 90 Fed. Reg. 35,393, 35,393–95 (July 28, 2025).

[122] Exec. Order No. 14,409, Promoting Advanced Artificial Intelligence Innovation and Security, § 3(a)–(c), 91 Fed. Reg. 34,565, 34,566 (June 5, 2026).

[123] *See* Anthropic, *Statement on the U.S. Government Directive to Suspend Access to Fable 5 and Mythos 5* (June 12, 2026), https://www.anthropic.com/news/fable-mythos-access; OpenAI, *Previewing GPT-5.6 Sol: A Next-Generation Model* (June 26, 2026), https://openai.com/index/previewing-gpt-5-6-sol/.

[124] NTIA Report, *supra* note 9, at 10–11, 40–47.

DRAFT

what is already available from other models.[125] Because highly capable Chinese open models already provide comparable capabilities in some critical domains,[126] U.S. open-weight releases appear, on the current record, to add little capability at the margin; restricting those releases would merely shift open-model development abroad and push users, at home and elsewhere, toward foreign substitutes without materially reducing overall risk.[127] Thus, the EAR's "published" exclusion should continue to apply by default for general-purpose models released with open weights.

In contrast, the marginal risk of publishing ITAR-controlled model weights would likely be higher. For model weights developed for use with defense articles, or weights that encode or reveal controlled technical data, comparable foreign substitutes will often not be available, so a release would supply adversaries with a defense-specific capability that they could not otherwise obtain. The benefits of publishing model weights, such as enhancing research or promoting competition, carry less force here. The existing regulatory structure aligns with that assessment. Unlike under the EAR, posting model weights on the internet generally would not remove them from the ITAR. The question for this category is therefore not whether to restrict publication but whether the existing restrictions can continue to withstand constitutional challenge.[128]

Although the current record does not support restricting general-purpose open releases, the marginal-risk analysis could change in the future. If a U.S. general-purpose model intended for open release were more advanced in relevant threat domains than foreign alternatives, temporarily restricting the release of its weights could meaningfully delay adversaries' access to the capabilities at issue. In such a case, the U.S. government's primary opportunity to intervene would be before the weights are published, while the developer still controls them.[129] The threshold question here is what form such a mandatory restraint could take. One option would be to create an exception to § 734.7(a)'s "published" exclusion for a defined category of model weights, together with an item-based licensing requirement that follows the regulatory structure used for machine-ready firearm files and specified encryption software.[130] Such a rule could adopt relevant capability thresholds to remain tailored to the risk, but to track the marginal-risk framework it would need to account for specific risk-relevant capabilities (not just a measure of overall capabilities), the availability of foreign substitutes, and changes in those capabilities and substitutes over time. Any item-based approach would also confront the Berman Amendment, which ECRA incorporates,[131] and First Amendment limits.

Another option would be an item-based specific notice of the kind used against Anthropic, which would be possible only if the weights remain subject to the EAR (i.e., before they are published). A third option would treat the release of the weights as U.S.-person

DRAFT

[125] *See supra* Part V.
[126] *See id.*
[127] *See id*.
[128] *See supra* Parts III.b.ii, III.d.
[129] *See supra* Parts III.a.ii–iii, IV.a, V.
[130] *See* 15 C.F.R. § 734.7(b)–(c); *id.* § 742.15(b).
[131] *See supra* Part III.d. Whether completed model weights qualify as protected "information or informational materials" remains uncertain, and the statutory exception for certain nonproliferation or antiterrorism controls would provide the government with a potential counterargument.

"support" for chemical- or biological-weapons development.[132] On that theory, which may be the government's strongest option, Commerce could take the position that releasing the weights is a service the developer knows may assist or benefit chemical- or biological-weapons development, if evaluations show that the model would provide substantial new capabilities in that area and that broad release creates an inherent risk of use for such development. Under § 744.6(b), a U.S. person with the requisite knowledge would need to seek a license, and BIS may also inform a person, individually or through a Federal Register notice, that a license is required for an activity that could involve covered "support."[133] If BIS found that a release of certain weights would make a material contribution to a covered end use, BIS would deny the license.[134] The "published" exclusion would not foreclose this scenario because § 744.6's restrictions on "support" are not limited to items subject to the EAR.

Although current law provides a plausible basis for this theory, it would be novel. The EAR has not previously treated the publication of information to the entire world as direct "support" for a weapons program. Section 744.6 applies to services that a U.S. person knows may assist or benefit a covered end use or end user—a connection that a release to the entire world may lack. A notice from BIS can impose a license obligation prospectively, but stretching "support" to reach such publication would test the limits of the regulation.

Moreover, regardless of whether the EAR reaches far enough to cover this approach, a license requirement for publication could plausibly qualify as a prior restraint that would present the First Amendment questions discussed in Part III.[135] The outcome of the constitutional analysis would, in part, turn on whether weights are determined to be protected expression, or, even if not, whether the restriction disproportionately burdens other protected First Amendment activities. The government's constitutional case would be strongest if evaluations showed that publication would materially facilitate the development of chemical- or biological-weapons in ways that could lead to mass casualties, and no comparable foreign substitute was available.[136] A sufficiently strong showing of such a mass-casualty risk could establish a compelling national security interest and justify a restriction narrowly drawn to address it, but the licensing regime would still require clear standards, appropriate limits, and prompt judicial review.

The government would have to weigh the litigation risk in this and any other factual scenario. A prior restraint on publication, particularly a broad one, could invite a constitutional challenge, and an adverse judicial decision could reach beyond model weights, limiting the government's authority to require approval before the release of unclassified technical data more

[132] *See* 15 C.F.R. § 744.6(b)(3).
[133] 15 C.F.R. § 744.6(c)(1).
[134] 15 C.F.R. § 744.6(e)(1).
[135] *See supra* Parts III.a.iv, III.d. The Berman Amendment could present a separate objection if applying the support restriction to publication operated as an indirect prohibition on the exchange of protected informational materials.
[136] The technological premise underlying that concern is not hypothetical, as researchers recently used AI systems trained on genetic data to design complete genetic blueprints for new viruses that infect bacteria, and some of those designs, when synthesized in a laboratory, produced functioning viruses substantially different from known natural viruses. *See* Samuel H. King et al., *Generative Design of Bacteriophages with Genome Language Models*, 393 Science eaec2657 (2026); *see also* Thomas V. Inglesby & Moritz S. Hanke, *AI-Designed Viral Genomes*, 393 Science 563 (2026). This study involved bacteriophages, which infect bacteria rather than human pathogens, so the work does not establish mass-casualty capability. But it does provide proof of concept for the generative design of complete, functional viral genomes.

DRAFT

generally.[137] Export controls, particularly the ITAR, have long relied on the ability to require pre-publication licensing of unclassified technical data, so litigation could jeopardize far more than the action at issue. Using a version of the June 2026 Anthropic export-control directive to block the release of model weights may stretch EAR authorities, operate as a prior restraint, and result in the type of crisis-driven improvisation that is hard to defend in litigation.

Bespoke authority governing the release of U.S. open models would be preferable, but it would be justified only if U.S. general-purpose open models intended for open release were beginning to provide threat-relevant capabilities beyond those available from foreign alternatives. In such a case, Congress should create a purpose-built authority to restrain the release of U.S. open-model weights in those narrow circumstances and provide the appropriate safeguards that a skeptical court's First Amendment review would demand (e.g., clear standards, appropriate limits, and prompt judicial review). A separate statute would avoid reliance on the ECRA (and IEEPA) authorities limited by the Berman Amendment, although any restraints would remain subject to the First Amendment. Federal law does not generally authorize the government to restrain the publication of privately created information, nor should it.[138] Where a narrow class of information raises concerns that might justify government-imposed restrictions, Congress can legislate to directly address that class. The TikTok experience reflects that pattern: after IEEPA-based restrictions were enjoined by the courts, Congress enacted a purpose-built statute, and the Supreme Court upheld it against First Amendment challenges.[139]

Until congressional action is needed to address risks from the release of U.S. open models, the government should pursue voluntary engagement with U.S. developers capable of a frontier-level open release. Voluntary efforts may not necessarily succeed in preventing releases from uncooperative developers, and a post-release export control directive, like that used with Anthropic, would have limited utility where the weights have already entered circulation. In that event, the government's remaining tools would be the same conduct-based authorities that apply to any published model, which reach outputs and downstream services rather than weights.

---

[137] *See Defense Distributed v. U.S. Dep't of State*, 838 F.3d 451, 461–76 (5th Cir. 2016) (Jones, J., dissenting) (concluding that the State Department's application of the ITAR's pre-publication review scheme to domestic online publication of lawful, unclassified technical data exceeded the AECA and violated the First Amendment, and warning that similar controls could extend to technical data concerning drones, cybersecurity, and robotic devices).

[138] Beyond export controls, other existing authorities reach adjacent conduct but do not cleanly reach publication. Section 337 of the Tariff Act of 1930 applies to "articles," and the International Trade Commission's jurisdiction under that provision does not extend to electronically transmitted digital data. *See ClearCorrect Operating, LLC v. Int'l Trade Comm'n*, 810 F.3d 1283, 1299 (Fed. Cir. 2015). Treasury's outbound-investment program regulates certain U.S. investment transactions involving covered Chinese AI developers, not the distribution or use of their models. *See* 31 C.F.R. §§ 850.209–.210, 850.217(d), 850.224(j)–(k). Federal procurement authorities, including exclusion and removal orders issued following recommendations by the Federal Acquisition Security Council, can exclude covered articles or sources from executive agency procurement and remove covered articles from executive agency information systems without reaching public availability. *See* 41 U.S.C. § 1323(c). And the Protecting Americans from Foreign Adversary Controlled Applications Act (PAFACA) reaches the hosting and distribution of foreign-adversary-controlled applications rather than stand-alone weight files. *See* Pub. L. No. 118-50, div. H, § 2(a), 138 Stat. 895, 955–56 (2024); *TikTok Inc. v. Garland*, 604 U.S. 56 (2025). The rare regimes that do restrain private publication are themselves narrow, purpose-built statutes, such as the Atomic Energy Act's Restricted Data provisions, 42 U.S.C. §§ 2014(y), 2274, and the Invention Secrecy Act, 35 U.S.C. §§ 181–188.

[139] *See TikTok Inc. v. Trump*, 507 F. Supp. 3d 92, 108–09 (D.D.C. 2020); *Marland v. Trump*, 498 F. Supp. 3d 624, 636–41 (E.D. Pa. 2020); Pub. L. No. 118-50, div. H, § 2(a), 138 Stat. 895, 955–56 (2024); *TikTok Inc. v. Garland*, 604 U.S. 56, 64–66, 80 (2025).

DRAFT

Similarly, Chinese developers abroad will generally remain beyond the direct reach of U.S. pre-release authorities and likely would not cooperate with any voluntary U.S. safety framework. To the extent that the U.S. government seeks to mitigate risks posed by foreign-origin open-weight models, its primary lever for doing so would be controls on upstream and downstream conduct with a U.S. nexus.

b. Applying the Marginal-Risk Framework to Outputs, Services, and Access

Existing authorities reach much of the relevant conduct involving model outputs, downstream services, and access to open models, and the marginal-risk framework should guide how agencies use them.

Particular model outputs may substantially advance the recipient's capabilities in a threat-relevant area beyond what realistic alternatives could provide. Other outputs may qualify as controlled information but would add little to the recipient's capabilities because similar information or assistance would be readily available from other sources. Treating model outputs as a general priority for enforcement would impose substantial compliance costs on those using open-weight models without materially reducing risks. Enforcement discretion should concentrate on high-consequence disclosures that materially advance the recipient's ability to carry out sensitive CBRN, cyber, or defense activity beyond what other sources of assistance could provide.[140]

Similarly, downstream conduct within reach of export-control authorities may provide capabilities in relevant domains that an adversary could not readily obtain elsewhere. Fine-tuning a model may enhance its capabilities in a sensitive domain or remove the model's safeguards. Hosting a model may provide reliable access to users who could not run the model locally. Integrating a model with proprietary data, tools, or infrastructure could help an adversary use it in ways that raise national security concerns. Conduct-based export controls are most useful when they target activities that materially enhance an adversary's capabilities.[141] BIS has already stated that U.S.-person services supporting AI-model training may require authorization when connected to covered WMD or military-intelligence activities.[142] It should analyze downstream services involving open weights in the same way and focus enforcement efforts on cases where services materially assist a covered WMD end use or military-intelligence end use or end user.

If such services proliferate and cannot be addressed case by case, BIS could issue a categorical rule, modeled on the semiconductor controls in § 744.6(c)(2)–(3), that would require licenses for specified services that could support a WMD program or military-intelligence end use or end user.[143] BIS would need to make factual findings comparable to those it made for semiconductors (e.g., that the services in question can contribute to covered end uses and that specific end uses are difficult to identify transaction by transaction due to China's military-civil fusion effort).[144] U.S. persons would need fair notice of what specific categories of conduct the

[140] *See supra* Parts III.e, V.
[141] *See supra* Parts III.a.iv, IV.b, V.
[142] *See supra* Part III.a.iv.
[143] *See id.*
[144] Implementation of Additional Export Controls: Certain Advanced Computing and Semiconductor Manufacturing Items; Supercomputer and Semiconductor End Use; Entity List Modification, 87 Fed. Reg. 62,186, 62,187 (Oct. 13,

DRAFT

rule covers, and BIS should expressly exclude published information and fundamental research, as it did for the semiconductor controls.[145]

Services involving Chinese open-weight models may warrant particular scrutiny. Agency guidance should identify red flags that should prompt enhanced due diligence before a company supplies a service and should treat violations involving those red flags as enforcement priorities. China's military-civil fusion policies can make military and civilian connections harder to assess, so any relevant military or intelligence ties could be a red flag. Moreover, evaluations have identified weak safeguards in some leading Chinese open-weight models.[146] Other red flags might therefore include the use of models with high-risk capabilities, degraded safeguards, and indications of deceptive or backdoor behavior, as well as requests to fine-tune or modify a model in ways that heighten risk. Red flags should inform diligence and enforcement priorities, though liability would still turn on the knowledge and end-use requirements.

These concerns have prompted reports of a possible "ban" on Chinese open-weight models.[147] A direct prohibition on U.S. persons downloading or locally using the weights would likely require new action under IEEPA or new legislation. The ICTS program is textually broad enough to reach such conduct, but its prior applications have generally regulated suppliers and commercial intermediaries rather than penalizing end users for continued possession or use.[148] A direct prohibition of this kind would also raise significant issues under the Berman Amendment and First Amendment, as described in Part III, and would be difficult to enforce after the weights had proliferated.

Instead, the government should focus its use of existing authorities at the points where it retains leverage. Upstream, it could slow Chinese labs' development of more capable models by limiting their access to frontier-scale compute, for which they still depend on U.S.- and allied-controlled sources. Where those inputs function as difficult-to-replace chokepoints, upstream measures can produce meaningful denial or delay. Downstream, the U.S. government could, where genuine risks are identified, limit U.S.-nexus services involving Chinese models. In addition to the EAR catch-all authorities described above, the ICTS program or a tailored IEEPA order could restrict U.S. hosting or inference involving a specified Chinese developer or platform.[149] These restrictions would be justified when the U.S. service would supply functionality or assistance not readily available elsewhere, or when it would create a demonstrated risk to U.S. systems. Under the Berman Amendment analysis discussed in Part III, restrictions stand on firmer ground when directed at the service relationship itself—supplying

DRAFT

2022) ("China's military-civil fusion effort makes it more difficult to tell which items are made for restricted end uses, thereby diminishing the effect of these existing controls.").

[145] *See* 15 C.F.R. § 744.6(d)(1)(ii).

[146] *See* Ctr. for AI Standards & Innovation, Nat'l Inst. of Standards & Tech., *CAISI Evaluation of DeepSeek AI Models Finds Shortcomings and Risks* (Sept. 30, 2025), https://www.nist.gov/news-events/news/2025/09/caisi-evaluation-deepseek-ai-models-finds-shortcomings-and-risks; Ctr. for AI Standards & Innovation, Nat'l Inst. of Standards & Tech., *CAISI Assessment of Z.ai's GLM-5.2* (July 17, 2026), https://www.nist.gov/news-events/news/2026/07/caisi-assessment-zais-glm-52.

[147] *See* Curi, *supra* note 5.

[148] *See supra* Part III.c; *see also* 15 C.F.R. § 791.2 (defining "ICTS Transaction" to include the acquisition, importation, transfer, installation, dealing in, or use of ICTS; defining "importation" to include electronic transmission; and defining a "Party" to include a buyer, purchaser, acquirer, and end user).

[149] *See supra* Part III.c.

inference to a platform, for example—than when their practical effect is to block the distribution of the weights.

An ICTS restriction on a hosted service or platform that maintains an ongoing connection to the foreign developer would resemble the program's prior applications. To reach transactions involving stand-alone weights that are no longer under the developer's control, however, Commerce would need to determine that the files themselves pose an undue or unacceptable risk. Weak safeguards or susceptibility to attack could support such a finding for transactions involving a particular model if the severity and likelihood of harm were sufficiently high. Repeated findings across a developer's models might support broader restrictions involving that developer, but it would be more difficult to justify a conclusion that applies across Chinese open-weight developers—or across developers linked to designated foreign adversaries more broadly.[150] A deliberately embedded "sleeper agent" or "backdoor" in a model with significant cyber or agentic capabilities would present a stronger basis for finding an undue or unacceptable risk because the risk would reside in the weights themselves even without an ongoing connection to the developer, thereby creating a risk to U.S. systems or infrastructure. Research shows that such behavior can be deliberately created and may persist through safety training,[151] but at present, available evaluations of current Chinese models identify only weak safeguards and susceptibility to attack, not backdoor behavior.[152] Evidence of a backdoor in one model could support restrictions on transactions involving that model, and evidence that the same developer used the practice across its models could support developer-wide restrictions. A broader prohibition covering transactions across multiple developers would require evidence of a systemic risk common to the defined class—e.g., a widespread or state-directed practice among developers in that class—and should be reserved for circumstances in which narrower restrictions or mitigation do not adequately address the risk. Thus, in appropriate circumstances, the ICTS program could serve as an effective tool addressing risks arising from downstream conduct with a U.S. nexus that involves Chinese or other foreign-adversary open-weight models.

Beyond the ICTS program, BIS and DDTC should issue guidance explaining how licensing and enforcement discretion will apply to model outputs and downstream services. They should clarify that a focus on marginal risk does not suspend existing prohibitions, which continue to apply according to their terms even where substitutes exist. Foreign availability may inform licensing policy or the allocation of enforcement resources, but it does not create an exception.[153] The guidance should also explain how agencies will distinguish outputs or services that add little to the recipient's capabilities beyond realistic alternatives from those that materially advance a sensitive activity. For model outputs in particular, providers may not be able to classify every output in real time, so the agencies should encourage risk-based compliance

DRAFT

---

[150] *See* 15 C.F.R. §§ 791.2, 791.103(c), 791.109(a) (providing for review and prohibition or mitigation of ICTS transactions or classes of transactions and identifying relevant factors for whether a Covered ICTS Transaction poses an undue or unacceptable risk, including technical capabilities, a discrete or persistent threat, severity of harm to critical infrastructure, likelihood of the threatened harm, and the availability of mitigation).

[151] *See* Evan Hubinger et al., *Sleeper Agents: Training Deceptive LLMs that Persist Through Safety Training* (Jan. 2024), https://arxiv.org/pdf/2401.05566 (constructing proof-of-concept examples in LLMs of backdoor behavior and finding that once a model exhibits deceptive behavior, standard techniques could fail to remove such deception and create a false impression of safety).

[152] *See* sources cited *supra* note 146.

[153] *See supra* Part V.

processes. The Center for AI Standards and Innovation (CAISI) should support this work by assessing the extent to which fine-tuning models, removing safeguards, and other downstream conduct could add capabilities beyond the base model and realistic alternatives, and by developing methods to test high-risk foreign-origin models for backdoor behavior.

c. Addressing Remaining Risks Through Domestic Law and International Cooperation

Export controls can, at most, address risks arising from transfers or services to foreign persons where a U.S. nexus exists. Risks arising from purely domestic conduct fall outside their reach and require other authorities. The design of those authorities—which may address misuse, liability, transparency, incident reporting, evaluation standards, and other mechanisms for mitigating the risks of open-weight models—lies beyond the scope of this paper. Whatever form those authorities take, research, commercial, and industrial incentives are likely to sustain a broad market for open-weight models. U.S. policy should therefore assume that such models will continue to be available and increasingly capable, and that the associated risks must be addressed through a coordinated set of tools in which export controls are only one component.

The United States and China also have a common interest in reducing the risk that open-weight models contribute to cross-border harms from cyber, biological, and other capabilities. In the near term, the countries could cooperate on limited technical topics such as best practices for assessing and reducing cybersecurity, CBRN, and loss-of-control risks.[154] The U.S.-China dialogue expected to begin in September 2026 could provide a forum for such cooperation.[155] Matt Sheehan, a senior fellow at the Carnegie Endowment, proposes a useful model of "AI safety in parallel," under which each country would strengthen its own safeguards, while fostering dialogue to enhance the safety of systems in both countries and increase mutual awareness of the other's safety practices.[156] Creating such information channels may help overcome the current lack of trust and establish a foundation for more ambitious future engagement.

While such cooperation should be prioritized, the United States should not weaken export controls (whether on advanced chips or on conduct) in exchange for Chinese promises on safety and security.[157] Advanced-chip controls can preserve a U.S. compute advantage, while targeted

---

[154] *See generally* Christina Knight & Scott Singer, *America and China Can Make AI Safer: Cooperation Is Necessary—and Possible*, Foreign Affs. (Apr. 7, 2026), https://www.foreignaffairs.com/united-states/america-and-china-can-make-ai-safer; Mark MacCarthy & Carl Schonander, *How the US and China Can Cooperate to Reduce Urgent AI Risks*, Brookings Inst. (July 14, 2026), https://www.brookings.edu/articles/how-the-us-and-china-can-cooperate-to-reduce-urgent-ai-risks/.

[155] *See* Laurie Chen, *U.S., China to Hold AI Talks in September, Sources Say*, Reuters (July 21, 2026), https://www.reuters.com/world/china/us-china-hold-ai-talks-september-sources-say-2026-07-21/.

[156] *See* Matt Sheehan, *A Path Forward on AI Safety for the United States and China*, Carnegie Endowment for Int'l Peace (July 23, 2026), https://carnegieendowment.org/emissary/2026/07/ai-safety-parallel-us-china.

[157] *Compare* Sebastian Mallaby, *I Went to China to See Its Progress on A.I. We Can't Beat It*, N.Y. Times (Apr. 13, 2026), https://www.nytimes.com/2026/04/13/opinion/china-ai-america-chipmakers.html (arguing that the United States should consider loosening chip controls to secure an AI-safety agreement), *with* Saif Khan et al., *Should the US Sell Hopper Chips to China?*, Inst. for Progress (Dec. 7, 2025), https://ifp.org/should-the-us-sell-hopper-chips-to-china/ (estimating that restrictions on advanced-chip exports can preserve a substantial U.S. compute advantage); *see also* Sheehan, *supra* note 156 (arguing that U.S.-China safety dialogue should not trade Chinese safety concessions for relaxed U.S. chip controls).

DRAFT

conduct-based controls prevent U.S. persons from providing services that materially support activities harmful to national security. The benefits of these controls are durable and should not be exchanged for promises that would be difficult to verify or enforce. The countries will remain strategic competitors, and export controls will continue to be necessary. Yet both countries have an independent interest in reducing cross-border AI risks, and sustained dialogue is an important first step.

VII. Conclusion

Open-weight models are likely to remain widely available, so U.S. policy must distinguish risks that export controls can reduce from those they cannot. Export controls are most useful when a reachable actor can be prevented from providing an adversary with capabilities not otherwise available, whether by releasing model weights, disclosing model outputs, or providing downstream services. Under the marginal-risk framework set out above, the EAR's "published" exclusion should be preserved as the default for general-purpose open-weight releases, and existing authorities should not be stretched to restrain the publication or download of model weights. If future U.S. open-weight releases begin to provide threat-relevant capabilities beyond foreign alternatives, Congress should provide purpose-built pre-publication authority before an urgent need arises, with clear standards, appropriate limits, and prompt judicial review. At present, agencies should focus their use of existing authorities on high-consequence model outputs and downstream activities that would materially advance an adversary's capabilities or create a demonstrated risk to U.S. systems. That is where the government's leverage endures. Once general-purpose weights are published and widely copied, item-based export controls lose their practical reach over the weights themselves, while fine-tuning, hosting, and related activities often still involve reachable actors. Risks beyond those chokepoints must be addressed through other domestic authorities and international cooperation.

DRAFT