How Supply-Side Constraints Reshape Demand-Side Innovation: Evidence from Chinese AI Firms' Open Source Activity

Yupeng Liu
University of Colorado Boulder
Email: yupeng.liu@colorado.edu

Tony W. Tong*
University of Colorado Boulder
Email: tony.tong@colorado.edu

Yuping Li
Singapore Management University
Email: yupingli@smu.edu.sg

This version: April 30, 2026

* Corresponding Author

Abstract: Firms in the Global South often depend on a small set of Global North incumbents for critical upstream technologies. How can they sustain innovation when access to these inputs is restricted not by any single partner, but by an across-the-board supply constraint? We develop the concept of dependence reconfiguration, arguing that constrained firms can redirect exchange relationships away from concentrated upstream suppliers toward distributed downstream users and developers. Open source software provides governance mechanisms that enable this shift: permissive licensing broadens downstream adoption, while open gatekeeping structures deepen external developer participation. We test this argument using a matched sample of Chinese software firms, comparing AI-related firms with non-AI firms before and after the October 2022 U.S. export controls on advanced computing hardware. We find that, after the shock, AI firms increased their release of permissively licensed open source projects and granted more external developers gatekeeping authority. These governance changes, in turn, increased community adoption and external contribution. The findings extend resource dependence theory by showing how firms manage upstream constraints through the cultivation of downstream relationships. We further identify open source innovation as a strategic pathway through which Global South firms sustain innovation under technological restrictions.

Keywords: Dependence reconfiguration, dependence structure, supply constraints, downstream innovation, open source innovation, open source governance.

Acknowledgement: The authors report no financial relationships or other potential conflicts of interest relevant to this research.

INTRODUCTION

Organizations in the Global South face a persistent structural challenge in technology-intensive industries. Advanced technologies, and the capabilities that underpin them, have long been concentrated among a small set of Global North incumbents, with emerging-market firms historically lagging on both fronts (Anand et al. 2021, Petricevic and Teece 2019). This concentration gives rise to a dependence structure characterized by high power imbalance and low mutual dependence, in which the dependent firm lacks bargaining leverage while the dominant party faces little pressure to accommodate it (Casciaro and Piskorski 2005). The asymmetry is further amplified when governments mobilize such supply dependencies as instruments of geopolitical competition, thereby reshaping the strategic options available to Global South firms in key sectors (Farrell and Newman 2019, Adarkwah et al. 2024, Han et al. 2024).

The Chinese AI industry exemplifies this form of concentrated upstream dependence. Since the introduction of the Transformer architecture, frontier AI development has operated under a compute-scaling paradigm, in which model performance improves as a predictable function of model size, dataset scale, and compute (Vaswani et al. 2017, Kaplan et al. 2020, Hoffmann et al. 2022). Notably, the advanced semiconductors that enable this training are designed and manufactured by a small set of U.S. and U.S.-allied firms, with NVIDIA holding a near-monopoly position in the market for AI training GPUs (*The New York Times* 2023, *Financial Times* 2023). Through 2022, Chinese AI firms invested heavily in the AI models and applications that depend on these chips, building their development capacity on an upstream supply base concentrated within U.S. regulatory jurisdiction. In October 2022, the U.S. Bureau of Industry and Security imposed export controls on these semiconductors, restricting the sale of advanced AI training chips, semiconductor fabrication equipment, and related technical services to Chinese entities, with the

stated aim of constraining China's ability to develop frontier artificial intelligence (Bureau of Industry and Security 2022, Allen 2022, Han et al. 2024, Luo and Van Assche 2023).

What can firms do under such supply-side dependence? Traditional resource dependence theory typically directs firms' response strategies toward the supply side of the exchange. For example, firms may acquire the supplier, turn to alternative suppliers, or exit the exchange (Pfeffer and Salancik 1978, Hillman et al. 2009, Jiang et al. 2023, Li et al. 2024). In the context of the U.S. export controls, however, none of these responses was available or feasible for Chinese AI firms. All potential suppliers of the restricted hardware fell under the same regulatory regime, cross-border acquisitions and alliances within the restricted domain were constrained by the same controls, and exiting the exchange would mean abandoning AI development altogether.

In this paper, we ask how firms reconfigure their exchange relationships to sustain innovation when upstream supply is constrained at the environmental level, rather than at the level of specific partners, as frequently studied in prior resource dependence research. We argue that firms can reconfigure dependence by shifting from *concentrated upstream* suppliers toward *distributed downstream* users, drawing on the demand-based view. It is well accepted that firm competitive advantage rests on demand-side users as well as on supply-side resources (Adner and Zemsky 2006, Priem 2007). When supply-side hardware is constrained, users remain a source of value creation through their adoption, adaptation, and further development of the firm's software (Baldwin and von Hippel 2011). Open source software provides an organizational form through which firms can engage users at scale. Under this model, source code is released publicly under licenses that permit external actors to use, modify, and redistribute it (Lerner and Tirole 2005, Nagle 2019). Firms typically host these projects in public repositories where users adopt the

technology and external developers submit proposed changes and contribute to its ongoing development.

How firms govern their open source projects has attracted increasing scholarly attention. Two governance choices, in particular, can significantly shape user engagement. First, firms can release code under permissive licenses (Lerner and Tirole 2005, Stewart et al. 2006), broadening the set of users able to adopt and build on the released code. Second, firms can grant external developers gatekeeping roles (O'Mahony and Karp 2022, Medappa et al. 2025), giving outside contributors a voice in shaping the project's direction. We argue that these governance choices can reconfigure the firm's exchange relationships away from concentrated upstream supply toward a distributed community of developers, thereby sustaining innovation through that community when the supply-side hardware becomes unavailable due to export controls.

We test this argument on a matched panel of Chinese software firms from October 2020 through October 2025, comparing Chinese AI-related software firms with Chinese non-AI software firms through a difference-in-differences design. The findings are consistent with the argument. After the October 2022 export controls, AI-related firms increased their release of permissively licensed open source projects and granted more external developers gatekeeping roles on their repositories. Users outside the firm adopted these projects at higher rates, and external developers contributed more to project development. These community-level responses operated through the two governance channels.

Our study makes two research contributions. First, we extend resource dependence theory by introducing dependence reconfiguration as a response to environmental-level resource constraints. Emerson's (1962, pp. 35–36) insight on "balancing operation" highlights the diffusion of dependence into "new relationships" as a strategy, yet existing research has primarily applied

his insight within the firm's existing upstream structure, treating it as substitution among similar exchange partners (Pfeffer and Salancik 1978, Hillman et al. 2009, Jiang et al. 2023). We argue that a firm's exchange relationships form a portfolio that includes not only the upstream suppliers on which its inputs depend, but also the *downstream* users who put the firm's software in use and extend its capabilities. When the upstream side of this portfolio is constrained across the board, firms can reconfigure their dependence away from concentrated upstream relationships toward distributed downstream relationships with a broad base of adopters and contributors. This shift brings resource dependence theory into conversation with the demand-based view (Adner and Zemsky 2006, Priem 2007) and research on user innovation (Baldwin and von Hippel 2011), thereby extending the theoretical scope of resource dependence theory. In essence, we advance existing theory by arguing that upstream dependence can be managed downstream, through governed relationships with firms' downstream user base.

Second, we contribute to research on organizations in the Global South by theoretically and empirically demonstrating open source governance as a strategic response to structural constraints. Prior work emphasizes how institutional voids and weak intermediary structures shape the strategies available to firms in these contexts (Khanna and Palepu 1997, Taussig and Delios 2015). We extend this perspective by showing how technology restrictions—constraints on access to critical upstream inputs—redirect firms' strategic behavior downstream toward open source innovation. We also identify a mechanism through which constrained firms sustain technology development: by releasing projects under open governance, firms can mobilize a global community of users and developers whose participation lies outside the scope of export restrictions (Zeitouni 2021). This distributed engagement substitutes, in part, for restricted upstream supplies, allowing firms to maintain innovation through downstream adoption and contribution. Moreover,

this strategy can enable firms to project influence beyond their boundaries by shaping technological trajectories. Through permissive licensing and more open gatekeeping, firms can expand adoption, attract complementary contributions, and establish *de facto* technical standards (Altman et al. 2022). Taken together, our findings advance a new view of Global South strategy: rather than merely adapting to supply constraints, firms can actively reconfigure their dependence in ways that convert upstream restrictions into downstream collaboration, thereby shaping technology development on a global scale.

THEORETICAL BACKGROUND

Resource Dependence and Upstream Supply Structure

Organizations depend on external actors for resources that are critical to their operations but not fully controlled internally. In resource dependence theory, the strength of such dependence is determined by two conditions: the importance of the focal resource to the firm and the availability of alternative sources from which that resource can be obtained (Emerson 1962, Pfeffer and Salancik 1978). When a resource is essential and alternatives are limited, firms become highly dependent on those actors that control access to it.

Upstream supply structure plays a central role in shaping these conditions. When supply is concentrated among a small set of providers, the lack of substitutable alternatives increases the focal firm's dependence on specific upstream actors. Two structural properties of such dependence determine what responses are available. First, power imbalance captures the asymmetry in how much each party needs the other, and second, mutual dependence captures the overall degree to which both parties rely on the exchange (Casciaro and Piskorski 2005). These properties have opposite effects on firms' ability to manage their dependence. When mutual dependence is high, both parties have incentives to restructure the relationship through mergers, alliances, or other

boundary-spanning arrangements (Hillman et al. 2009). When power imbalance is high, the stronger party faces no such incentive, and the weaker party lacks the bargaining power to initiate change. Buyer-supplier relationships illustrate this distinction. Mutual dependence between buyers and suppliers can improve buyers' performance by enabling joint action and richer information exchange, whereas dependence asymmetry favoring suppliers will erode buyers' performance (Gulati and Sytch 2007). Supply structure thus shapes not only the level of dependence but also whether that dependence can be managed through interaction within the existing exchange relationship.

Prior research highlights a range of strategies through which firms manage dependence on upstream suppliers, including diversifying the supplier base, forming alliances, vertical integration, and political action (Pfeffer and Salancik 1978, Hillman et al. 2009, Jiang et al. 2023). These strategies share an implicit assumption that the firm retains some room to maneuver within its supply environment, whether by switching to an alternative source, acquiring a supplier, or bringing a critical activity in-house. When access to a critical upstream resource is constrained not at the level of individual exchange relationships but across the set of upstream actors that control the resource, the above options become substantially more difficult to implement. The problem firms face is no longer how to improve their position within the existing supply structure, but the broader question of how to sustain value creation when that structure cannot be readily managed.

Environmental-Level Restriction on Upstream Dependence

Resource dependence theory distinguishes between two levels at which dependence operates (Pfeffer and Salancik 1978). At the partner level, dependence arises from a specific exchange relationship and can be addressed by altering the terms of that relationship or by seeking alternative partners within the same supply domain. At the environmental level, dependence arises from how

a critical resource is distributed across the entire field of potential suppliers. Switching partners does not resolve asymmetry, because the constraint lies in the supply environment itself rather than in any particular exchange relationship (Jiang et al. 2023, Xia 2011). Thus, when dependence operates at the environmental level, the conventional strategies of supplier diversification and alliance formation may no longer be effective.

Firms facing environmental-level supply constraints can shift toward internal capability building as a means of managing upstream dependence. Under such conditions, firms can increase R&D investment to mitigate supply-side constraints (Liu et al. 2025) or redirect R&D toward technologically complementary domains (Helfat 1997). These responses operate within organizational boundaries, focusing on firms' internal adjustments to address upstream dependence.

Such internal responses, however, may be insufficient, especially for firms in the Global South. In technology-intensive industries, critical upstream resources remain concentrated among Global North actors, placing Global South firms in a position of persistent dependence (Anand et al. 2021, Petricevic and Teece 2019). The resulting dependence configuration is marked by high power imbalance and low mutual dependence, in which the dependent firm lacks bargaining power and the dominant party faces no compulsion to accommodate (Casciaro and Piskorski 2005). This imbalance is reinforced by several structural conditions underlying the relationship between the Global North and Global South. Decades of path dependence in the global distribution of technological capabilities (Awate et al. 2015), compounded by institutional voids that constrain indigenous capability development (Khanna and Palepu 1997), sustain the asymmetry between upstream suppliers and downstream adopters. Cross-border barriers further restrict the feasible set of responses by raising the costs of mergers, alliances, and other boundary-spanning strategies that

rely on regulatory openness across jurisdictions (Luo and Van Assche 2023, Xia et al. 2014). For Global South firms, the internal response strategies described above face additional structural disadvantages that make their limits more binding.

Dependence Reconfiguration and Opportunities for Contribution

Resource dependence theory identifies three classes of dependence management strategy: absorption, countervailing, and avoidance (Pfeffer and Salancik 1978, Hillman et al. 2009, Jiang et al. 2023). Absorbing strategies expand organizational boundaries through acquisitions, alliances, or vertical integration. Countervailing strategies cultivate alternative sources to reduce dependence on any single provider. Avoiding strategies withdraw from or diversify away from the constrained exchange. Under environmental-level constraints, these paths encounter serious limitations. Absorption requires access to targets that are themselves subject to the constraint; when regulatory boundaries harden, cross-border internalization loses viability. Countervailing fails because all sources within the supply domain fall under the same restriction, the condition Emerson (1962, p. 35) characterizes as dependent party being "denied alternative sources for achieving those goals". Avoidance is not viable when asset positions tie the firm to the constrained domain and exit would mean abandoning past investment. The classic "balancing operation" strategy in Emerson's (1962) framework, which emphasizes extending the network by cultivating additional exchange partners of the same type, also runs in the upstream direction and is constrained for the same reason. The conventional set of strategies thus shares a directional limitation, because they all seek to address dependence within the upstream supply structure.

Yet the upstream supply structure is not the only direction in which dependence relationships can be restructured. Restricting what the firm can obtain from *upstream* suppliers does not limit how *downstream* actors can engage with the firm, and firms possessing redeployable

knowledge resources can redirect them toward alternative pathways when established paths become unavailable. Rather than continuing to search for solutions within the upstream supply structure, then, firms can build new exchange relationships in the downstream direction, with actors who engage with and build upon the firm's existing resources. Emerson (1962, p. 37) predicts that actors in unbalanced dependence relations can restore equilibrium through "the diffusion of dependency into *new relations* in a network." Drawing on this inspiration, prior applications have directed this diffusion upstream, toward additional suppliers of the restricted resource (Xia 2011); by contrast, we redirect the diffusion to downstream. Specifically, we argue that when upstream network extension is constrained, firms can diffuse dependence into relationships with distributed communities of downstream actors. The resulting dependence structure differs from the upstream configuration along two dimensions: upstream dependence is concentrated among few suppliers, with high power imbalance and low mutual dependence; by contrast, downstream dependence is distributed across many actors, with lower power imbalance and higher mutual dependence.

This shift constitutes a form of what we term dependence reconfiguration. This idea extends existing strategies of dependence management beyond absorbing, countervailing, and avoiding by identifying a downstream pathway in which the firm moves from *concentrated upstream dependence* on a small number of powerful suppliers to *distributed downstream dependence* across a broad set of exchange partners. In doing so, it addresses a gap in resource dependence theory, which has predominantly theorized responses within the existing upstream supply structure and has given less attention to strategies that redirect dependence in a new direction—downstream, in our case. The specific mechanisms through which this reconfiguration

operates, and the governance choices that shape downstream relationships, are the focus of our hypotheses development below.

HYPOTHESES DEVELOPMENT

Hardware-Software Complementarity and the Demand-Based View

In theory, activities are complementary when doing more of one increases the returns to doing more of another (Milgrom and Roberts 1990). In the software industry, hardware and software exhibit strong complementarity: advanced hardware enables more capable software, and sophisticated software extracts the full performance potential of the hardware on which it runs (Shapiro and Varian 1999). The value a firm derives from one layer of its technology stack therefore depends on the configuration of adjacent layers. Upstream constraints on components and downstream relationships with users affect the focal firm's position differently (Adner and Kapoor 2010). When the upstream hardware complement is constrained, the channel through which software previously realized value through joint deployment with that hardware is restricted. The software the firm owns retains potential value, but that value now depends less on supply-side bundling with scarce hardware, but more on the relationship between software and the users who adopt and extend it.

The demand-based view of competitive advantage provides the foundation for this alternative path. At its core, this view suggests that competitive advantage depends not only on a firm's supply-side resources but also on how effectively the firm connects with heterogeneous customers on the demand side (Adner and Zemsky 2006, Schmidt et al. 2024). Importantly, supply-side and demand-side sources of advantage can be distinct; when the former is restricted, value creation may still proceed through the latter (Priem 2007). Consistent with this argument, recent research finds that when demand conditions change, firms redirect their innovation search toward

opportunities presented by the altered demand environment (Zhong et al. 2021). We complement this stream of research by developing a set of hypotheses on whether and when firms redirect their innovation activity to the demand side as a result of exogenous constraints on the supply side. Drawing on the demand-based view, our core argument is that when the supply of hardware complement is constrained, firms shift strategic emphasis toward the demand side of the technology stack, where software value arises not through joint deployment with the upstream complement but through the engagement of users and developers in the developer community who adopt, adapt, and build upon it.

User Innovation and Open Source Software

In the software industry, users account for a disproportionate share of functionally novel improvements, particularly when their needs are highly heterogeneous (Franke and von Hippel 2003, Lakhani and von Hippel 2003). Internal development and external knowledge acquisition are complementary, with each amplifying the value of the other (Cassiman and Veugelers 2006). When supply constraints limit what the firm can achieve through internal effort alone, the relative value of contributions from users increases. Engaging users as participants in the innovation process becomes important for sustaining competitive advantage under constraints (Lakhani and von Hippel 2003), as demand-side value creation through user engagement does not depend on the restricted supply-side channel (Adner and Kapoor 2010).

Open source software, defined as software whose source code is publicly available and can be used, modified, and redistributed under specified license terms (Lerner and Tirole 2005), provides the organizational form through which firms engage users at scale. In theory, firms grant outside actors access to their code, opening markets for complementary innovation around the firm's technology (Lakhani and von Hippel 2003). In practice, a firm releases an open source

project by publishing a repository on an online platform such as GitHub under an open source license, thereby opening the project to outside developers and organizations who can freely adopt it, adapt it to their own purposes, and contribute to its further development (Lin and Maruping 2022, Conti et al. 2025). Each release opens a new channel through which the firm engages distributed downstream actors, and the number of projects a firm releases provides an observable measure of the scale at which it pursues such engagement.

When supply constraints restrict firms' access to the upstream hardware complement, the supply-side channel through which software previously realized value is curtailed. The demand-side source of advantage—sustained through user adoption, adaptation, and contribution—remains available and provides a viable alternative channel. Specifically, by releasing open source projects under open source licenses, supply-constrained firms can build exchange relationships with distributed downstream actors, including users, developers, and contributors. This process deepens the firm's dependencies on these distributed downstream parties, thereby shifting from concentrated upstream dependence toward the distributed downstream dependence described above. Open source project release therefore provides a way supply-side dependence can be reconfigured in the software industry, leading to the following baseline hypothesis:

Baseline Hypothesis: *Following supply-side constraints on access to upstream hardware complement, software firms will increase open source project release.*

This hypothesis establishes the baseline prediction that constrained firms will release more open source projects. However, whether such releases enable firms to engage a broad and active downstream community depends on the governance choices they make, particularly the types of licenses adopted to distribute the projects and the kinds of developers deployed to gatekeep community contributions, which we turn to in the hypotheses below.

Effects of Supply-Side Constraints on Open Source Project Governance

Building on the baseline hypothesis, we now develop a set of hypotheses on firms' governance choices in releasing open source projects, which can shape downstream community engagement. For open source projects to mobilize the distributed community, developers across the open source community must adopt the code, adapt it to their own use, and contribute their work back to the project. We argue that how effectively a firm mobilizes such participation depends on the governance mechanisms it deploys. In open source projects, governance refers to the rules and structures that determine what outside actors can do with the code and how much influence they have over a project's development (O'Mahony and Karp 2022). Consistent with prior work, we focus on two dimensions of this governance. First, permissive licensing shapes the breadth of participation by determining the conditions under which the project can be adopted and redistributed. Second, gatekeeping openness affects the depth of participation by determining who can shape the project's technical trajectory.

Adoption of Permissive Licenses

Software licenses define the conditions under which an open source project may be used, modified, and redistributed (Lerner and Tirole 2002), and license decisions are central to how open source communities engage with such projects (He et al. 2020). One key distinction across licenses is the extent to which they restrict downstream use, ranging from permissive to restrictive. Permissive licenses, such as the MIT, BSD, and Apache families, impose only minimal attribution requirements, allowing downstream users to incorporate the code into proprietary products or redistribute modified versions under commercial terms. By contrast, restrictive licenses, most notably those in the GPL family, require that any modified versions downstream remain under the same open license. This reciprocity requirement effectively deters firms that develop and sell

proprietary software, because incorporating GPL-licensed code would obligate them to release the resulting product under the same open terms (Open Source Initiative n.d.). Research has found that projects adopting permissive licenses tend to attract greater downstream user interest than those adopting restrictive licenses (Lerner and Tirole 2005, Stewart et al. 2006).

The difference between permissive and restrictive licensing is fundamentally a difference in how broad a user base the released project can reach. Restrictive licenses exclude commercial and related users whose business models require proprietary distribution, thereby narrowing the downstream user base to those willing to accept the same open terms. Permissive licenses admit these users alongside all others, producing a wider and more diverse set of adopters.

A broader user base benefits the releasing firm in two ways. First, a more heterogeneous user base includes adopters working with the technology across a wider range of environments. Each of these users holds context-specific information about performance and adaptation needs that is costly for the releasing firm to acquire on its own (Franke and von Hippel 2003, Lakhani and von Hippel 2003). Second, wider adoption generates network effects (Shapiro and Varian 1999). Each adopter who integrates the technology into their own products and workflows becomes a potential source of improvements that feed back into the project; the more adopters a project has, the faster it improves and the more attractive it becomes to subsequent adopters (von Hippel and von Krogh 2003). As adoption grows, the technology can become a *de facto* standard, and firms can profit from widely adopted open source technologies through complementary services and proprietary extensions built on the open technology base (West 2003, Casadesus-Masanell and Llanes 2011). Both advantages grow with the breadth of the downstream user base and are therefore greater under permissive than under restrictive licensing.

When supply-side constraints restrict firms' access to the upstream hardware complement, firms will depend more heavily on contributions from downstream users. The scale and diversity of these contributions depend on the breadth of the downstream user base, which in turn depends on how permissive the license is. Permissive licensing widens this user base and, in doing so, expands the set of downstream relationships through which the firm reconfigures its dependence (Emerson 1962, Casciaro and Piskorski 2005). Firms under constraints therefore have stronger incentives to release their open source projects under permissive licenses.

Hypothesis 1a: *Following supply-side constraints on access to upstream hardware complement, software firms will increase their release of open source projects under permissive licenses.*

Deployment of External Gatekeepers

While permissive licensing determines how broadly a project can be adopted, gatekeeping openness determines how deeply external developers can participate in shaping the project's development. In open source projects, this participation is governed through write access to the source code, which grants the ability to merge proposed changes into the main codebase, review contributions from others, and extend these privileges to additional developers (O'Mahony and Ferraro 2007, Medappa et al. 2025). Write access divides contributors into two groups. Those with write access act as gatekeepers, defining the project's trajectory and deciding which contributions enter the codebase (Gousios et al. 2014). Those without write access can propose changes, report issues, and submit work for review, but cannot modify the codebase directly. A firm can retain write access with its own employees or extend it to external developers who have earned trust through sustained contribution.

When a firm extends gatekeeper roles to external developers, it shares control over the project's development with actors outside the firm. This governance choice motivates external contribution through two channels. First, extending gatekeeper roles to non-employees signals that the project operates under a more democratic model and is open to outside participation (O'Mahony and Ferraro 2007, O'Mahony and Karp 2022). When external developers already assume gatekeeper roles, potential contributors can see that outside contributions have a credible pathway to shape the project's direction, thereby increasing their willingness to invest sustained effort. Second, participation in open source governance functions as a labor market signal, and this signal strengthens with the level of governance responsibility. Developers who attain higher-ranked governance positions in open source projects have greater employability (Hann et al. 2013, Xu et al. 2020). When a firm opens gatekeeper roles to external developers, it creates a pathway through which contributors can attain positions that carry both technical authority over the project and visible labor market value. While extending gatekeeping could introduce coordination costs, firms that value sustained external contributions also have stronger incentives to accept these costs.

When supply-side constraints restrict firms' access to the upstream hardware complement, what the firm can achieve through internal development alone becomes more limited, making it depend more heavily on sustained contributions from external developers. Wider adoption expands the downstream user base, yet the firm also relies on developers who contribute improvements and extend the technology in ways its internal team does not. Extending gatekeeper roles to these developers transforms a unilateral relationship into a distributed form of engagement. This, in turn, creates the mutual dependence between the firm and its downstream contributors that characterizes the depth of the reconfigured dependence structure (Gulati and Sytch 2007). Firms under constraints therefore have stronger incentives to open their gatekeeper roles to external developers.

Hypothesis 1b: *Following supply-side constraints on access to upstream hardware complement, software firms will increase the use of external developers in gatekeeping their open source projects.*

Effects of Governance Change on Community Engagement

The preceding hypotheses argue that supply-side constraints lead firms to increase their release of projects under permissive licenses and extend gatekeeper roles to external developers. Next, we proceed to examine whether these strategies lead to increased engagement with the open source projects from the developer community. Community engagement with an open source project takes two forms, which we refer to as community adoption and community contribution.

Community adoption occurs when an external developer or organization integrates the firm's technology into their own products or workflows. On platforms such as GitHub, adoption often takes the form of forking, in which an external actor creates a copy of the project's repository that can then be developed along its own trajectory (Conti et al. 2025). Forking extends the reach of the technology across organizational boundaries and creates a new downstream exchange relationship in which the adopter depends on the project's codebase. On the other hand, community contribution occurs when an external developer proposes improvements to the original project by submitting a pull request, which the project's maintainers can review and, if accepted, integrate into the codebase (Gousios et al. 2014, Medappa et al. 2025). Contribution represents a deeper form of engagement than adoption, because the contributor invests effort in shaping the project's trajectory, enabling the firm to benefit from development work beyond its own internal effort. In our dependence reconfiguration framework, community adoption expands the breadth of new downstream exchange relationships, while community contribution deepens these relationships into the mutual dependence between the firm and its downstream partners.

We draw on the private-collective model of innovation to develop hypotheses on the relationship between open source project governance and community engagement (von Hippel and von Krogh 2003). Developers participate in open source projects when the private returns to participation exceed the associated costs. These private returns include the ability to use and modify the technology for one's own purposes (Lerner and Tirole 2002), learning through contributing to collectively developed software that translates into productivity gains (Nagle 2018), career signaling through visible contributions and governance roles that translate into labor market premiums (Hann et al. 2013, Xu et al. 2020), and access to external knowledge for technology development through engagement with actively maintained open source communities (Conti et al. 2025).

Permissive licensing and gatekeeping openness each shape these private returns and, in turn, affect both community adoption and community contribution. For community adoption, permissive licensing relaxes legal constraints that would otherwise prevent commercial users from integrating the technology into proprietary products, while an open, democratic governance model signals that the project will evolve in response to community needs and is therefore a more friendly and reliable technology to adopt (O'Mahony and Ferraro 2007, O'Mahony and Karp 2022). For community contribution, permissive licensing enlarges the pool of users from which contributors can emerge, while extending gatekeeping roles to external developers offers a visible governance position that functions as a career signal and demonstrates that contributions from outside the firm are valued and can meaningfully shape the project's development trajectory (Hann et al. 2013). The governance strategies adopted by firms under supply-side constraints proposed in the preceding two hypotheses will therefore lead to increases in both community adoption and community contribution. In the dependence reconfiguration framework, these community

responses constitute evidence on firms' shift from concentrated upstream dependence to distributed downstream dependence. Thus, we hypothesize:

Hypothesis 2a: *Following supply-side constraints on access to upstream hardware complement, software firms' governance changes in open source projects will facilitate community adoption.*

Hypothesis 2b: *Following supply-side constraints on access to upstream hardware complement, software firms' governance changes in open source projects will facilitate community contribution.*

Figure 1 summarizes the structure of these hypotheses: Export controls (explained in the empirical context section below) operate as the supply-side shock that drives the two governance responses (H1a and H1b), each of which, in turn, mediates the effect of the shock on community adoption (H2a) and community contribution (H2b).

----------Insert Figure 1 about Here----------

DATA AND METHODS

Empirical Context

The advanced semiconductors that power AI development are among the most concentrated upstream resources in global technology production, and the recent restrictions on Chinese firms' access to them provide the empirical setting for our study. On October 7, 2022, two months after the passage of the CHIPS and Science Act (CHIPS and Science Act of 2022, Luo and Van Assche 2023), the United States Bureau of Industry and Security issued export controls on the advanced computing chips used for AI training, on advanced semiconductor fabrication equipment, and on U.S. persons' services in support of Chinese chip development, together with an extension of the Foreign Direct Product Rule covering advanced computing (Bureau of Industry and Security 2022,

Allen 2022). The controls applied uniformly to all Chinese firms whose products depend on the restricted hardware. The October 2022 rule provides a quasi-natural experiment that constrained Chinese AI firms' access to upstream hardware, marking the discrete shift at which the critical inputs on which these firms depended became systematically unavailable.

Frontier AI development depends on advanced computing hardware in ways that are not easily substitutable. Under the compute-scaling paradigm that has driven AI progress since the introduction of the Transformer architecture, model performance improves as a predictable function of compute, such that training frontier models requires large quantities of advanced semiconductors (Vaswani et al. 2017, Kaplan et al. 2020, Hoffmann et al. 2022). As of 2022, Chinese domestic chip manufacturers were not able to produce semiconductors at the performance levels required for frontier AI training (Allen 2022), and export controls restricted access to the U.S.-designed chips on which Chinese firms had built their development infrastructure. Because hardware and software in AI are complementary (Moore et al. 2018, Xu et al. 2010), these controls aimed to constrain China's frontier AI development by restricting critical hardware inputs (Allen 2022).

Since late 2022, Chinese firms have released a succession of widely adopted open source AI models, frameworks, and developer tools. In fact, the share of globally adopted open-weight AI models originating in China has grown substantially, reaching levels comparable to those of U.S. firms by 2024 (Maslej et al. 2025). Our empirical focus is on AI firms' open source project release and governance decisions on GitHub, the dominant public platform for firm-level open source activity (Malgonde et al. 2023, Conti et al. 2025). Specifically, firms post their open source work as public repositories that they create and maintain, and they govern each repository through the open source license they attach, which specifies the terms under which external actors may use

and redistribute the code. Users and developers outside the firm engage with these repositories through publicly observable actions: they fork a repository to adopt the technology for their own use, and they submit pull requests that propose changes to the codebase. These actions map directly to the dependent variables in our hypotheses. Permissive project release and gatekeeping openness capture the two governance dimensions predicted by H1a and H1b, while forks and external pull requests capture the community adoption and contribution outcomes predicted by H2a and H2b. In our design, we treat the October 2022 export controls as a quasi-natural experiment on Chinese AI firms' upstream hardware supply and use a matched-panel difference-in-differences design to test whether firms under this constraint reconfigure their dependence through open source governance.

Data and Sample

To construct our dataset, we combine company-level data on Chinese software firms and their funding histories from Crunchbase, monthly records of firms' open source activity drawn from the GitHub event archive and the GitHub REST API, and individual-level employment histories from Revelio Labs. We follow a growing body of firm-level research that measures open source engagement using GitHub event data aggregated to the firm-month (Nagle 2019, Wright et al. 2023, Conti et al. 2025), and we use the Revelio Labs employment records to distinguish internal from external developers in constructing our gatekeeping-openness measure.

Our base sample comprises Chinese software firms tracked in Crunchbase. We focus on software firms for two reasons. First, they are the most active participants on GitHub, where open source release and community engagement are directly observable. Second, they offer cleaner identification than Chinese hardware and semiconductor firms, which face ambiguous incentives under the export controls because the same restrictions that limit access to foreign chips may

simultaneously increase domestic demand for their products as substitutes. With software firms, only those whose products depend on advanced AI computing hardware are directly exposed to the controls, yielding a treated group and a control group drawn from the same industry.

We restrict the sample to firms headquartered in China whose Crunchbase industry categories include "Software." We match Crunchbase records to GitHub organizations by company names and domains, and retain firms for which an identifiable GitHub organization account can be verified. The resulting panel tracks these firms at a monthly frequency from October 2020 through October 2025, yielding a five-year observation window spanning twenty-four months before and thirty-six months after the October 2022 policy.

We classify each firm along a five-level AI involvement rubric using a large language model (see Appendix A for details), following recent work that applies LLMs to annotate text at scale (Gilardi et al. 2023). The five levels range from non-AI (level 1, firms with no meaningful AI component) to AI feature (level 2, firms that incorporate AI functionality such as a chat-bot into otherwise non-AI products, but whose core operations do not depend directly on the restricted computing hardware), AI application (level 3, firms whose products are centered on AI), AI infrastructure (level 4), and foundation-model developer (level 5). The classifier takes as input each firm's name, full business description, and Crunchbase industry categories. We exclude five firms whose 2022 Crunchbase profiles do not flag AI as a primary activity but whose 2025 profiles document substantive AI activity. We further restrict the sample to firms founded before October 2022, excluding 17 firms whose post-shock entry cannot reflect a strategic response to the export controls.

The treated group consists of 143 Chinese software firms classified at the AI application level or higher. These firms build products centered on artificial intelligence, including AI

applications, AI infrastructure, and foundation models, all of which depend on the advanced computing hardware restricted by the October 2022 export controls. The control group consists of 521 Chinese software firms classified at the AI feature or non-AI levels (levels 2 or 1, respectively). To validate the LLM-based classification, we manually inspected the model's outputs against the underlying Crunchbase text and corrected labels where warranted. In the main analysis, we operationalize group membership as a binary indicator, and we also examine a continuous measure of AI involvement in a supplementary analysis presented later.

We match treated and control firms using Coarsened Exact Matching (Iacus et al. 2012). Because treated and control firms may differ on observable characteristics that independently shape open source activity, we balance the two groups *ex ante* so that the difference-in-differences estimates are not driven by compositional differences. The matching balances on firm age (quartile-binned), GitHub organization age (quartile-binned), venture-capital backing, public listing status, and an indicator equal to one for medium or large firms (those with fifty-one or more employees), following the OECD micro-and-small-enterprise threshold (OECD 2024). The matched sample for analysis comprises 664 firms and 36,337 firm-month observations.

Table 1 reports covariate balance before and after matching. The three binary covariates balance well post-match, with standardized mean differences (SMDs) all at or below 0.14. Firm age and GitHub organization age retain higher SMDs because CEM enforces exact balance on coarsened quartile bins rather than on continuous values, and treated firms are a structurally younger cohort of Chinese AI entrants than the broader Chinese software population. Firm fixed effects in our main specification absorb differences across firms, and we assess the parallel-trends assumption directly through an event-study analysis.

----------Insert Table 1 about Here----------

Variables

We use five dependent variables that correspond to the predictions developed in the preceding section. *New repository creation* measures the number of new open source projects the firm releases, corresponding to the baseline hypothesis. The next two, *permissive project release* and *gatekeeping openness*, capture the firm's governance over its open source projects. The final two, *forks* and *external pull requests*, represent community adoption and community contribution, respectively.

New repository creation is the firm's monthly count of newly created, non-fork repositories. We count only original repositories, excluding those forked from other accounts, because a forked repository is a copy of externally developed code rather than an original release. *New repository creation* therefore measures the rate at which the firm releases new open source projects.

Permissive project release is the firm's monthly count of newly created, non-fork repositories released under permissive licenses. We classify repositories as permissive or restrictive following the Open Source Initiative's canonical registry (Open Source Initiative n.d., Stewart et al. 2006). Permissive licenses, such as MIT, BSD, and Apache, allow broad downstream use, modification, and commercialization, whereas restrictive licenses in the GPL family require that derivatives remain open, thereby limiting commercial reuse and broad adoption (Lerner and Tirole 2005). We aggregate this monthly count as a trailing six-month sum to smooth short-run noise in the release series. H1a predicts that firms under supply-side constraints increase their release of permissively licensed projects.

Gatekeeping openness is an indicator for firm-months in which external developers make up the majority of maintainers holding write access to the firm's active open source projects. Write access grants permission to merge changes directly into a project's codebase, making write-access

holders the effective decision-makers over the project's direction. We identify write-access holders from GitHub push events and classify each contributor as internal or external by matching their GitHub profile to employment histories from Revelio Labs. A contributor is considered active from the month of their first push through three months after their last observed push, accommodating typical activity gaps without treating long-time contributors as departed after a single inactive month. When external contributors account for more than half of active write-access holders, the majority of those able to modify the firm's codebase are outsiders rather than employees, signaling more democratic project governance (O'Mahony and Ferraro 2007, O'Mahony and Karp 2022). H1b predicts that firms under supply-side constraints shift toward this more open governance strategy.[1]

Forks is the monthly count of instances in which external developers copy one of the firm's repositories into their own account to use, modify, or build on independently. Forks are a standard measure of community adoption used in prior research based on GitHub data, as each fork records a downstream actor's decision to take up the firm's code (Nagle 2019, Conti et al. 2025). H2a predicts that forks increase when the firm's governance encourages broader downstream adoption through the release of open source projects under permissive licenses or the use of external developers in gatekeeping their open source projects.

External pull requests is the firm's monthly count of pull requests submitted by developers who are not employees of the focal firm. A pull request is a proposed change to the codebase, typically a bug fix, performance improvement, or new feature, that the firm reviews and accepts or rejects. We focus on external pull requests to isolate engagement originating outside the firm,

[1] GitHub's event archive records write-access decisions only when contributors exercise their permissions, so firms whose projects are inactive or maintained by a single developer do not generate observable gatekeeping records. The regressions involving *gatekeeping openness* and the corresponding mediation analyses are restricted to the subsample of 386 firms (11,923 firm-month observations) for which these decisions are observable.

as these contributions represent effort and innovation the firm would otherwise need to generate internally (O'Mahony and Ferraro 2007). H2b predicts that external pull requests increase when the firm's governance encourages broader developer contribution.

Our analysis includes three time-varying, firm-level controls—venture-capital backing, public listing status, and prior acquisition history—which capture access to capital, public-market exposure, and ownership status. Time-invariant firm characteristics are absorbed by firm fixed effects.

Table 2 presents summary statistics for the matched sample, separately for the treated and control groups.

----------Insert Table 2 about Here----------

Empirical Approach

To test our hypotheses, we estimate a difference-in-differences specification with two-way fixed effects:

$$Y_{it} = \beta \cdot (\text{Treated}_i \times \text{Post}_t) + \alpha_i + \delta_t + \gamma' X_{it} + \varepsilon_{it},$$

where Y_{it} is one of the five dependent variables for firm i in month t. The firm fixed effect α_i absorbs time-invariant firm characteristics, including location, founding cohort, and core technology stack. The month fixed effect δ_t absorbs aggregate shocks common to all Chinese software firms in a given month, including broader open source trends and macroeconomic conditions. X_{it} is the set of three time-varying firm-level controls described above, and ε_{it} is the error term. The coefficient of interest, β, captures the average effect of the October 2022 policy shock on treated firms relative to control firms. Standard errors are clustered at the firm level.

The design relies on the assumption that, absent the policy, treated and control firms would have followed parallel trajectories in the dependent variables. The October 2022 export controls

applied uniformly to all Chinese firms whose products depend on advanced AI computing hardware, supporting the interpretation of β as the causal effect of the shock on the treated group. We assess the parallel-trajectories assumption directly through an event-study specification presented in the Results section, and we address global open source trends through the cross-country comparison with Indian firms in a placebo test reported later.

For H2, we test whether the policy's effects on community engagement operate through two mechanisms: permissive licensing and gatekeeping openness. We follow the Baron and Kenny (1986) mediation framework. The indirect effect of the policy through a given mediator is defined as the product of the policy's effect on the mediator and the mediator's effect on the dependent variable, controlling for treatment. We test each mediator separately for each of the two H2 dependent variables, forks and external pull requests, yielding four single-mediator tests. For inference, we report both the Sobel test and bias-corrected 95 percent confidence intervals from a cluster-resampled bootstrap of 5,000 replications (Preacher and Hayes 2008). By resampling entire firms rather than individual firm-month observations, the bootstrap accounts for the clustered panel structure of our data.

RESULTS

Main Effects

Table 3 reports the difference-in-differences estimates of the export controls' effect on the five dependent variables, with standard errors clustered at the firm level (in parentheses). Column 1 tests the baseline hypothesis that supply-side constraints induce treated firms to release more open source projects. The estimate on Treated × Post is 0.0306 ($p = 0.010$), indicating that treated firms increased their rate of new repository release by about three percent relative to the control group, consistent with the baseline prediction.

----------Insert Table 3 about Here----------

Columns 2 and 3 test H1a and H1b, which predict that affected firms respond by increasing the release of permissively licensed projects and by adopting more open gatekeeping of those projects. Treated firms' permissive project release rose by 0.0709 ($p = 0.005$), corresponding to roughly a seven-percent increase relative to the control group. In addition, the probability that external developers hold majority write access rose by 15.7 percentage points ($p = 0.005$). Taken together, the two estimates indicate that treated firms increased the use of permissive licensing and admitted more external developers into gatekeeping roles, consistent with H1a and H1b.

Columns 4 and 5 report the total effects of the export controls on community adoption, measured by forks of the firm's repositories, and on community contribution, measured by external pull requests submitted by outside developers. Treated firms received an 11 percent increase in forks ($p = 0.022$) and a 10 percent increase in external pull requests ($p = 0.023$). To assess whether these increases operate through the governance changes suggested in H1a and H1b, we turn to the mediation mechanisms predicted by H2 in the next subsection.

Mediating Mechanisms

H2a and H2b predict that the two forms of governance change, permissive licensing and gatekeeping openness, mediate the export controls' effects on community adoption (H2a) and community contribution (H2b). Table 4 reports the estimated indirect effects for each mediator across both dependent variables, along with Sobel tests and bias-corrected 95 percent confidence intervals from bootstrapping.

----------Insert Table 4 about Here----------

The indirect effects provide support for both H2a and H2b. Permissive licensing and gatekeeping openness each exhibit a statistically significant indirect effect on *forks* (0.0229 and

0.0330, respectively) and on *external pull requests* (0.0401 and 0.0559, respectively), with all bias-corrected 95 percent confidence intervals excluding zero (Table 4). The two governance mechanisms mediate between 10.3 and 39.6 percent of the total effect on each dependent variable, with permissive licensing accounting for the larger share of the total effect mediated in both cases.

Taken together, the mediation results support the central claim of the dependence-reconfiguration argument. The increased use of permissive licensing and the admission of external developers into gatekeeping roles represent the channels through which firms under supply-side constraints shift from concentrated upstream dependence toward more distributed downstream dependence.

Event-Study Analysis

Figure 2 plots event-study coefficients for four focal dependent variables, estimated using the approach of Borusyak et al. (2024), which replaces the binary post indicator with a sequence of leads and lags around October 2022. In the pre-shock period, the coefficients cluster near zero with overlapping confidence intervals across all four panels. Joint *F*-tests of the pre-shock coefficients yield p-values above the 0.10 threshold for each dependent variable (reported in the notes under Figure 2), consistent with the parallel-trends assumption underlying our main specification.

----------Insert Figure 2 about Here----------

In the post-shock period, the coefficients diverge from zero in the predicted direction within the first quarter after October 2022, and this divergence persists through the end of the observation window. The visual pattern aligns with the regression estimates in Table 3, indicating that the observed effects arise in response to the export controls rather than continuing pre-existing trends.

Robustness Checks

We test the robustness of our main estimates to three methodological choices: the matching procedure, the measures of the two governance variables, and the treatment classification. We address each in turn.

First, we re-estimate the main specification under three alternative matching procedures to ensure that the results are not sensitive to our CEM implementation. Each alternative varies the coarsening by using finer or coarser bin definitions for firm age, GitHub organization age, and firm size. Across the four focal dependent variables (*permissive project release*, *gatekeeping openness*, *forks*, and *external pull requests*), the coefficients on *Treated* × *Post* retain their sign and statistical significance under each alternative specification, and each falls within the 95 percent confidence interval of the main estimate (Table 5). These results indicate that the findings are not driven by specific matching procedures.

----------Insert Table 5 about Here----------

Second, we vary the measurement of the two governance variables. For *gatekeeping openness*, we re-estimate the models using the underlying continuous share of external write-access holders in place of the binary majority indicator, and we vary the length of the active window across zero, three, and six months around the baseline. For *permissive project release*, we re-estimate the models under trailing windows ranging from three to twelve months around the baseline six-month window. Across all alternative specifications, the estimated effects on the governance variables retain their sign, remain comparable in magnitude to the main estimates, and remain statistically significant (results not tabulated to conserve space). These results indicate that the governance findings reflect underlying behavioral shifts rather than specific operationalization choices.

Finally, we vary the classification used to define the binary treatment. The main specification separates firms at the AI application level or higher (level 3 or above) from those at the AI feature level or below (level 2 or 1)—a discrete contrast that captures the first-order comparison but does not exploit variation in firms' positions along the AI value chain. We therefore re-estimate the model using the five-level ordinal classification that underlies the binary treatment, coded from 1 (non-AI firms) to 5 (foundation-model developers), interacted with the post-shock indicator. We find that the estimated effect increases monotonically with AI exposure across the ordinal scale for each dependent variable, and the interaction term is statistically significant as well (results not tabulated to conserve space). These results are consistent with the argument that firms positioned higher up the AI value chain are more directly exposed to the export controls.

Placebo Tests

The October 2022 export controls coincided with the launch of ChatGPT on November 30, 2022 (Maslej et al. 2025), which triggered a worldwide expansion of AI-related activity and could have independently shifted open source release and community engagement across all AI firms. To separate the policy effect from this contemporaneous global shock, we re-estimate the main specification on a comparable panel of Indian software firms. Because India was not subject to the US export controls, Indian AI firms experienced the ChatGPT expansion without the hardware restriction that bound Chinese firms. We construct the Indian panel using the same Crunchbase-to-GitHub matching procedure. Panel A of Table 6 reports the difference-in-differences estimates for the Indian sample across the four focal dependent variables featured in H1 and H2. All four coefficients are statistically indistinguishable from zero, ruling out the ChatGPT-era expansion in AI activity as the driver of the effects estimated in the Chinese sample.

----------Insert Table 6 about Here----------

The India placebo test indicates that Indian AI firms did not respond to the October 2022 export controls. A triple-difference specification extends this analysis by estimating how much larger the response of Chinese AI firms was relative to Indian AI firms over the same period. We pool the Chinese and Indian samples and estimate a three-way interaction of AI treatment, the post-shock indicator, and the China indicator, together with firm and month fixed effects and the controls used throughout. Panel B of Table 6 reports the coefficient on the triple interaction term *Treated* × *Post* × *China*, which captures the extent to which the change for Chinese AI firms (relative to non-AI firms) exceeds the corresponding change for Indian firms during the post-shock period. The coefficient is positive and statistically significant at the 5 percent level for three of the four dependent variables (*permissive project release*, *forks*, and *external pull requests*) and marginally significant at the ten-percent level for *gatekeeping openness*. The triple-difference results thus reinforce the paper's central claim that upstream supply-side constraints drive firms to reconfigure their dependence toward more distributed, downstream open source engagement, above and beyond any contemporaneous AI-sector trends.

DISCUSSION AND CONCLUSION

This paper asks how Global South firms reconfigure their exchange relationships to sustain innovation when upstream supply is constrained at the environmental level. We study this question in the setting of Chinese AI firms after the October 2022 U.S. export controls on advanced computing hardware. Results from a matched panel of Chinese software firms show an interesting pattern. After the export controls, Chinese AI firms increased their release of permissively licensed open source projects and granted more external developers gatekeeping roles on their repositories. Users outside the firm adopted these projects at higher rates, and external developers contributed more to project development. Importantly, both community responses operated through the two

governance shifts. These results are robust to various tests including a placebo test comparing Chinese AI firms with Indian AI firms, which are not subject to the export controls. Below, we discuss the implications of these findings for resource dependence theory and for research on organizations in the Global South.

Dependence Reconfiguration and Resource Dependence Theory

Resource dependence theory has long argued that firms can manage dependence through strategies such as absorption, countervailing, and avoidance, all of which focus on reducing dependence within the upstream supply structure (Pfeffer and Salancik 1978, Hillman et al. 2009, Jiang et al. 2023). These strategies share an implicit assumption that firms retain some discretion over their upstream relationships. While powerful in many settings, these strategies have less value when constraints arise at the environmental level—conditions that are increasingly salient under techno-geopolitical rivalry and cross-border decoupling (Farrell and Newman 2019, Petricevic and Teece 2019, Luo and Van Assche 2023, Witt et al. 2023)—where restrictions apply broadly across upstream actors and limit the scope for firm-level actions in the upstream.

In this setting, we argue and show that firms can still manage dependence by treating their exchange relationships as a portfolio that spans both upstream and downstream actors. Emerson's (1962, pp. 35–36) insight on "balancing operation" describes the diffusion of dependency into new relationships as a means to restore balance under unbalanced dependence. We argue that this diffusion can run downstream as well as upstream, toward distributed users and developers who adopt, adapt, and build on the firm's output. We term this response *dependence reconfiguration*. It extends resource dependence theory by identifying a downstream response in which firms manage concentrated upstream dependence by redirecting effort to building relationships with distributed users and developers. The extension brings resource dependence theory into direct conversation

with two literatures. The demand-based view treats customers as a source of competitive advantage and identifies the demand side as a distinct locus of value creation (Adner and Zemsky 2006, Priem 2007, Schmidt et al. 2024). Research on user innovation establishes users and developer communities as productive participants in the innovation process (Lakhani and von Hippel 2003, Baldwin and von Hippel 2011). The concept of dependence reconfiguration links the two literatures to resource dependence theory by specifying when firms shift to the demand side as a strategic response to environmental supply constraints, and by treating user innovation as a managed response that firms shape through deliberate governance choices.

Our results sharpen this argument by showing how the upstream and downstream sides of the firm's exchange-relationship portfolio differ in their structural properties (Casciaro and Piskorski 2005, Gulati and Sytch 2007). Upstream dependence in our setting is concentrated among a small number of suppliers and exhibits high power imbalance and low mutual dependence; by contrast, downstream dependence is distributed across many users and developers and exhibits lower power imbalance and higher mutual dependence. Two governance choices shape this downstream structure. Permissive licensing widens the user base by removing reciprocity requirements that would deter commercial users, and gatekeeping openness deepens engagement by giving external developers authority over the project's technical trajectory. Theoretically, it's worth noting that gatekeeping openness helps operationalize the mutual-dependence mechanism that prior research associates with stronger relational performance (Gulati and Sytch 2007).

Open Source Innovation as a Strategic Response for Global South Firms

Our findings extend research on organizations in the Global South. Prior work has largely focused on how firms in these contexts navigate domestic institutional voids and weak intermediary structures (Khanna and Palepu 1997, Taussig and Delios 2015), and, in the face of supply-side

constraints on critical upstream inputs, has emphasized input substitution or internal capability building. In this study, we identify a distinct response that operates beyond firm boundaries. Specifically, we show that Chinese AI firms facing severe supply constraints sustain technological development by releasing open source projects and making deliberate governance choices that mobilize a global community of users and developers in the downstream.

It is well established that hardware and software are complementary inputs in AI development (Milgrom and Roberts 1990, Brynjolfsson et al. 2021). The value of joint development of the two layers of the AI stack means that software-side innovation can compensate, at least partially, for hardware-side constraints. This complementarity, combined with the widespread open-sourcing of AI models and applications, highlights the strategic role of open source innovation for Global South firms facing restricted access to critical inputs. Specifically, we see two mutually-reinforcing functions that open source serves for these firms. First, as a novel organizational form, open source innovation sits outside the scope of semiconductor export restrictions (Zeitouni 2021), allowing firms to engage the global community of users and developers when restrictions on hardware bind. Open source thus offers a technology development channel through which firms mobilize the community to advance the software side of the AI stack. DeepSeek's recent release of its V4 model illustrates this mechanism in operation. Developed under continuing U.S. restrictions on advanced chips, the model achieves frontier-competitive performance through architectural innovations on the software side that substantially reduce the hardware demand for inference and training. Its technical report explicitly emphasizes "software-hardware co-design" with chip vendors (DeepSeek-AI 2026, p. 16), demonstrating our logic in practice. In addition, the V4 model was released publicly under permissive terms and received

rapid global adoption (Maslej et al. 2025), mobilizing downstream community engagement in line with our theoretical argument and findings.

Second, open source release can strengthen a firm's position in global technological influence. Broad diffusion of technology through open source attracts complementary contributions (Lakhani and von Hippel 2003) and helps establish *de facto* technical standards (West 2003), positioning the firm to influence the trajectory of technical development in its domain (Altman et al. 2022). Reflecting these dynamics, even firms with proprietary origins increasingly shift core development onto open foundations (Haese and Peukert 2025). Through the open source channel, constraints on upstream input access can be strategically reconfigured into downstream influence: technologies controlled by the firm can become embedded in development trajectories that extend beyond domestic markets. This enables Global South firms to participate in shaping the direction of global technological evolution even when their access to critical upstream resources remains restricted.

Boundary Conditions and Future Directions

We acknowledge several limitations due to the scope of the project. First, the open source mechanism we identify is highly relevant to the software industry but may be less applicable to other settings. The Chinese AI industry provides a well-suited context to test our theory because semiconductor export restrictions apply only to hardware, and because hardware and software are highly complementary to each other, especially in AI development. In industries where firms' output is physical and tightly coupled to the restricted upstream input, the open source mechanism will be less applicable, and firms facing environmental-level supply constraints would need to develop other response strategies to manage their dependence.

A second limitation concerns upstream substitution. Our data capture downstream open source release and governance choices but do not allow us to examine whether Chinese AI firms adopt domestically available, lower-performance chips as substitutes for restricted U.S. hardware. Whether firms combine downstream community engagement with upstream substitution, or instead face a trade-off between the two, remains an open question. Addressing this question would require fine-grained data on hardware sourcing and procurement decisions, which we leave to future research.

A natural direction for future research concerns the *direction* of firms' innovation under sustained constraints. Constrained firms may redirect their innovative efforts away from compute-intensive foundation model development, where upstream restrictions bind most tightly, toward application-oriented product development, where dependence on advanced chips is less acute and released models can be deployed and adapted across domains, as DeepSeek's development trajectory illustrates. Linking governance choices and community engagement to firms' innovation direction, and ultimately performance outcomes, will sharpen our understanding of the theoretical and practical value of the dependence reconfiguration strategy we propose for Global South firms, which are likely to remain dependent on a concentrated set of Global North incumbents for critical technologies in the foreseeable future.

REFERENCES

Adarkwah GK, Dorobantu S, Sabel CA, Zilja F (2024) Geopolitical volatility and subsidiary investments. *Strategic Management Journal* 45(11):2275–2306.

Adner R, Zemsky P (2006) A demand-based perspective on sustainable competitive advantage. *Strategic Management Journal* 27(3):215–239.

Adner R, Kapoor R (2010) Value creation in innovation ecosystems: How the structure of technological interdependence affects firm performance in new technology generations. *Strategic Management Journal* 31(3):306–333.

Allen GC (2022, October 11) *Choking off China's access to the future of AI*. Center for Strategic and International Studies. https://www.csis.org/analysis/choking-chinas-access-future-ai

Altman EJ, Nagle F, Tushman ML (2022) The translucent hand of managed ecosystems: Engaging communities for value creation and capture. *Academy of Management Annals* 16(1):70–101.

Anand J, McDermott G, Mudambi R, Narula R (2021) Innovation in and from emerging economies: New insights and lessons for international business research. *Journal of International Business Studies* 52(4):545–559.

Awate S, Larsen MM, Mudambi R (2015) Accessing vs. sourcing knowledge: A comparative study of R&D internationalization between emerging and advanced economy firms. *Journal of International Business Studies* 46(1):63–86.

Baldwin C, von Hippel E (2011) Modeling a paradigm shift: From producer innovation to user and open collaborative innovation. *Organization Science* 22(6):1399–1417.

Baron RM, Kenny DA (1986) The moderator-mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology* 51(6):1173–1182.

Borusyak K, Jaravel X, Spiess J (2024) Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies* 91(6):3253–3285.

Brynjolfsson E, Rock D, Syverson C (2021) The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics* 13(1):333–372.

Bureau of Industry and Security (2022, October 13) Implementation of additional export controls: Certain advanced computing and semiconductor manufacturing items; Supercomputer and semiconductor end-use; Entity List modification. *Federal Register* 87(197):62186–62215. https://www.govinfo.gov/content/pkg/FR-2022-10-13/pdf/2022-21658.pdf

Casadesus-Masanell R, Llanes G (2011) Mixed source. *Management Science* 57(7):1212–1230.

Casciaro T, Piskorski MJ (2005) Power imbalance, mutual dependence, and constraint absorption: A closer look at resource dependence theory. *Administrative Science Quarterly* 50(2):167–199.

Cassiman B, Veugelers R (2006) In search of complementarity in innovation strategy: Internal R&D and external knowledge acquisition. *Management Science* 52(1):68–82.

CHIPS and Science Act of 2022, Pub. L. No. 117-167, § 103, 136 Stat. 1366 (2022). https://www.congress.gov/117/plaws/publ167/PLAW-117publ167.pdf

Conti A, Peukert C, Roche MP (2025) Beefing IT up for your investor? Engagement with open source communities, innovation, and startup funding: Evidence from GitHub. *Organization Science* 36(4):1551–1573.

DeepSeek-AI (2026) *DeepSeek-V4: Towards highly efficient million-token context intelligence* [Technical report]. https://huggingface.co/deepseek-ai/DeepSeek-V4-Pro/blob/main/DeepSeek_V4.pdf

Emerson RM (1962) Power-dependence relations. *American Sociological Review* 27(1):31–41.

Farrell H, Newman AL (2019) Weaponized interdependence: How global economic networks shape state coercion. *International Security* 44(1):42–79.

Financial Times (2023) China′s internet giants order $5bn of Nvidia chips to power AI ambitions (August 9). https://www.ft.com/content/9dfee156-4870-4ca4-b67d-bb5a285d855c

Franke N, von Hippel E (2003) Satisfying heterogeneous user needs via innovation toolkits: The case of Apache security software. *Research Policy* 32(7):1199–1215.

Gilardi F, Alizadeh M, Kubli M (2023) ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences* 120(30): e2305016120.

Gousios G, Pinzger M, van Deursen A (2014) An exploratory study of the pull-based software development model. In *Proceedings of the 36th International Conference on Software Engineering (ICSE '14)* 345–355. ACM.

Gulati R, Sytch M (2007) Dependence asymmetry and joint dependence in interorganizational relationships: Effects of embeddedness on a manufacturer's performance in procurement relationships. *Administrative Science Quarterly* 52(1):32–69.

Haese J, Peukert C (2025) Open at the core: Moving from proprietary technology to building a product on open source software. *Management Science*. Advance online publication.

Han P, Jiang W, Mei D (2024) Mapping U.S.–China technology decoupling: Policies, innovation, and firm performance. *Management Science* 70(12):8386–8413.

Hann I-H, Roberts JA, Slaughter SA (2013) All are not equal: An examination of the economic returns to different forms of participation in open source software communities. *Information Systems Research* 24(3):520–538.

He VF, Puranam P, Shrestha YR, von Krogh G (2020) Resolving governance disputes in communities: A study of software license decisions. *Strategic Management Journal* 41(10):1837–1868.

Helfat CE (1997) Know-how and asset complementarity and dynamic capability accumulation: The case of R&D. *Strategic Management Journal* 18(5):339–360.

Hillman AJ, Withers MC, Collins BJ (2009) Resource dependence theory: A review. *Journal of Management* 35(6):1404–1427.

Hoffmann J, Borgeaud S, Mensch A, Buchatskaya E, Cai T, Rutherford E, de Las Casas D, Hendricks LA, Welbl J, Clark A, Hennigan T, Noland E, Millican K, van den Driessche G, Damoc B, Guy A, Osindero S, Simonyan K, Elsen E, Rae JW, Vinyals O, Sifre L (2022) Training compute-optimal large language models. In *Advances in Neural Information Processing Systems 35 (NeurIPS 2022)*.

Iacus SM, King G, Porro G (2012) Causal inference without balance checking: Coarsened exact matching. *Political Analysis* 20(1):1–24.

Jiang H, Luo Y, Xia J, Hitt MA, Shen J (2023) Resource dependence theory in international business: Progress and prospects. *Global Strategy Journal* 13(1):3–57.

Kaplan J, McCandlish S, Henighan T, Brown TB, Chess B, Child R, Gray S, Radford A, Wu J, Amodei D (2020) Scaling laws for neural language models. *arXiv Preprint arXiv:2001.08361*.

Khanna T, Palepu K (1997) Why focused strategies may be wrong for emerging markets. *Harvard Business Review* 75(4):41–51.

Lakhani KR, von Hippel E (2003) How open source software works: "Free" user-to-user assistance. *Research Policy* 32(6):923–943.
Lerner J, Tirole J (2002) Some simple economics of open source. *Journal of Industrial Economics* 50(2):197–234.
Lerner J, Tirole J (2005) The economics of technology sharing: Open source and beyond. *Journal of Economic Perspectives* 19(2):99–120.
Li J, Shapiro DM, Ufimtseva A, Zhang P (2024) Techno-nationalism and cross-border acquisitions in an age of geopolitical rivalry. *Journal of International Business Studies* 55(9):1190-1203.
Lin Y-K, Maruping LM (2022) Open source collaboration in digital entrepreneurship. *Organization Science* 33(1):212–230.
Liu X, Liu Y, Makarin A, Wen Y (2025) *Export controls and innovation in sanctioned countries* (Working Paper No. 25–004). Harvard Business School.
Luo Y, Van Assche A (2023) The rise of techno-geopolitical uncertainty: Implications of the United States CHIPS and Science Act. *Journal of International Business Studies* 54(8):1423–1440.
Malgonde OS, Saldanha TJV, Mithas S (2023) Resilience in the open source software community: How pandemic and unemployment shocks influence contributions to others' and one's own projects. *MIS Quarterly* 47(1):361–390.
Maslej N, Fattorini L, Perrault R, Gil Y, Parli V, Kariuki N, Capstick E, et al. (2025, April) *The AI Index 2025 Annual Report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University.
Medappa PK, Srivastava SC, Favaron SD (2025) Write access provisioning and organizational ownership in open source software projects: Exploring the impact on project novelty and survival. *Research Policy* 54(8):105284.
Milgrom P, Roberts J (1990) The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review* 80(3):511–528.
Moore AW, Hebert M, Shaneman S (2018) The AI stack: A blueprint for developing and deploying artificial intelligence. *Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR IX*, vol. 10635 (SPIE, Bellingham, WA), 45–54.
Nagle F (2018) Learning by contributing: Gaining competitive advantage through contribution to crowdsourced public goods. *Organization Science* 29(4):569–587.
Nagle F (2019) Open source software and firm productivity. *Management Science* 65(3):1191–1215.
OECD (2024) *Enterprises by business size* [Indicator]. OECD Data. Retrieved April 22, 2026, from https://doi.org/10.1787/31d5eeaf-en
O'Mahony S, Ferraro F (2007) The emergence of governance in an open source community. *Academy of Management Journal* 50(5):1079–1106.
O'Mahony S, Karp R (2022) From proprietary to collective governance: How do platform participation strategies evolve? *Strategic Management Journal* 43(3):530–562.
Open Source Initiative. (n.d.). *Open source licenses by category*. Retrieved April 21, 2026, from https://opensource.org/licenses-old/category
Petricevic O, Teece DJ (2019) The structural reshaping of globalization: Implications for strategic sectors, profiting from innovation, and the multinational enterprise. *Journal of International Business Studies* 50(9):1487–1512.

Pfeffer J, Salancik GR (1978) *The External Control of Organizations: A Resource Dependence Perspective* (Harper & Row, New York).
Preacher KJ, Hayes AF (2008) Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods* 40(3):879–891.
Priem RL (2007) A consumer perspective on value creation. *Academy of Management Review* 32(1):219–235.
Schmidt J, Priem RL, Zanella P (2024) Customers, markets, and five archetypical value creation logics: A review of demand-side research in strategic management. *Journal of Management* 50(6):2309–2342.
Shapiro C, Varian HR (1999) *Information Rules: A Strategic Guide to the Network Economy* (Harvard Business School Press, Boston).
Stewart KJ, Ammeter AP, Maruping LM (2006) Impacts of license choice and organizational sponsorship on user interest and development activity in open source software projects. *Information Systems Research* 17(2):126–144.
Taussig M, Delios A (2015) Unbundling the effects of institutions on firm resources: The contingent value of being local in emerging economy private equity. *Strategic Management Journal* 36(12):1845–1865.
The New York Times (2023) The desperate hunt for the AI boom's most indispensable prize (August 16). https://www.nytimes.com/2023/08/16/technology/ai-gpu-chips-shortage.html
Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I (2017) Attention is all you need. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)* 5998–6008.
von Hippel E, von Krogh G (2003) Open source software and the "private-collective" innovation model: Issues for organization science. *Organization Science* 14(2):209–223.
West J (2003) How open is open enough? Melding proprietary and open source platform strategies. *Research Policy* 32(7):1259–1285.
Witt MA, Lewin AY, Li PP, Gaur A (2023) Decoupling in international business: Evidence, drivers, impact, and implications for IB research. *Journal of World Business* 58(1):101399.
Wright NL, Nagle F, Greenstein S (2023) Open source software and global entrepreneurship. *Research Policy* 52(9):104846.
Xia J (2011) Mutual dependence, partner substitutability, and repeated partnership: The survival of cross-border alliances. *Strategic Management Journal* 32(3):229–253.
Xia J, Ma X, Lu JW, Yiu DW (2014) Outward foreign direct investment by emerging market firms: A resource dependence logic. *Strategic Management Journal* 35(9):1343–1363.
Xu L, Nian T, Cabral L (2020) What makes geeks tick? A study of stack overflow careers. *Management Science* 66(2):587–604.
Xu X, Venkatesh V, Tam KY, Hong S-J (2010) Model of migration and use of platforms: Role of hierarchy, current generation, and complementarities in consumer settings. *Management Science* 56(8):1304–1323.
Zeitouni S (2021) Milling the f/loss: Export controls, free and open source software, and the regulatory future of the internet. *New York University Journal of Legislation and Public Policy* 23(3):905–952.

Zhong W, Ma Z, Tong TW, Zhang Y, Xie L (2021) Customer concentration, executive attention, and firm search behavior. *Academy of Management Journal* 64(5):1625–1647.

Table 1. Covariate balance, pre- and post-matching

	Pre-matching			Post-matching		
Covariate	Treated	Control	SMD	Treated	Control	SMD
Firm age (months)	59.1	99.5	0.613†	65.0	87.6	0.363†
GitHub organization age (months)	17.0	33.6	0.446†	18.6	31.1	0.336†
Venture-capital backed	0.610	0.585	0.050	0.685	0.714	0.063
Publicly listed	0.031	0.091	0.249	0.035	0.065	0.139
Medium or large firm ($\geq$ 51 employees)	0.610	0.561	0.098	0.604	0.603	0.003
n treated		159			143	
n control		685			521	

Notes. Covariate means for treated and control firms before and after Coarsened Exact Matching (Iacus, King, & Porro, 2012), with standardized mean differences (SMDs; Austin, 2011). The pre-matching sample is all Chinese software firms in Crunchbase with an identifiable GitHub account and a founding date before October 2022 ($n = 844$; five contaminated controls removed). The post-matching sample is the 664-firm analytic panel ($n = 664$; an additional 17 post-treatment-founded firms removed). Firm age and GitHub organization age are reported in months; remaining covariates are binary proportions. Binary covariates are matched exactly in CEM; continuous covariates are coarsened to quartile bins before matching and reported here in original units. After matching, venture-capital backing and the medium-or-large firm indicator achieve SMDs below the strict 0.10 benchmark, and publicly listed (0.139) sits within the 0.10–0.25 acceptance range (Austin, 2011). SMDs marked † remain elevated because CEM guarantees exact balance on quartile bins rather than on continuous values.

Table 2. Summary statistics by treatment status (matched analytic sample)

	Control ($n = 521$)		Treated ($n = 143$)	
	Mean	SD	Mean	SD
Dependent variables				
New repository creation (count)	0.101	0.586	0.088	0.553
Permissive project release (count)	0.040	0.330	0.041	0.288
Gatekeeping openness	0.698	0.459	0.711	0.454
Forks (count)	2.390	15.822	1.019	6.665
External pull requests (count)	4.812	41.293	0.923	6.381
Firm controls				
Venture-capital backed	0.794	0.405	0.828	0.377
Publicly listed	0.068	0.252	0.050	0.218
Previously acquired	0.035	0.184	0.007	0.084
Medium or large firm	0.604	0.489	0.604	0.489
Firm age (months)	118.1	73.5	95.4	55.1
GitHub organization age (months)	61.3	42.9	48.5	39.6

Notes. Statistics computed at the firm-month level over the observation window (October 2020 through October 2025). The treated group consists of 143 Chinese software firms whose AI value-chain classification places them at the AI application level or higher; the control group consists of 521 Chinese software firms classified at the AI feature or non-AI levels. Count dependent variables are reported in their original (unlogged) units; the regressions in Table 3 apply log-of-one-plus transformations to all four count-based dependent variables. *Permissive project release (count)* is the monthly count of newly released repositories under permissive licenses; the regressions use $\ln(1 + \text{six-month rolling sum})$ of this monthly flow, which smooths short-run noise. The other three count-based dependent variables enter the regressions as $\ln(1 + y_{it})$ of the monthly count. *Gatekeeping openness* is an indicator equal to one when external write-access holders hold the majority of project gatekeeping. Among the firm controls, *medium or large firm* is an indicator for firms with fifty-one or more employees, separating firms above the OECD micro-and-small-enterprise threshold (OECD, 2024); the four control variables (venture-capital backing, public listing, prior acquisition, medium or large firm) are time-invariant at the firm level, while firm age and GitHub organization age advance with the panel.

Table 3. Effect of the export controls on open-source activity, governance, and community engagement

	OSS activity	*OSS governance*		*Community engagement*	
	(1) New repository creation (Baseline)	(2) Permissive project release (H1a)	(3) Gatekeeping openness (H1b)	(4) Forks (H2a)	(5) External pull requests (H2b)
Treated × Post	0.0306***	0.0709***	0.1568***	0.1056**	0.1013**
	(0.0118)	(0.0250)	(0.0554)	(0.0461)	(0.0446)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes
Calendar-month FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Firms (clusters)	664	664	386	664	664
Observations	36,337	36,337	11,923	36,337	36,337
R^2	0.297	0.485	0.670	0.771	0.724

Notes. Each column reports a separate difference-in-differences regression with firm and calendar-month fixed effects. The treated group consists of 143 Chinese software firms classified at the AI application level or higher; the control group consists of 521 Chinese software firms classified at lower levels. Treated and control firms are matched using Coarsened Exact Matching (Iacus, King, & Porro, 2012). Post equals one for months from October 2022 onward. The dependent variable in Column 1 (Baseline) is the log of one plus the monthly count of new non-fork repositories. Column 2 (H1a) is the log of one plus the rolling six-month count of newly released repositories under permissive licenses (MIT, BSD, Apache, and similar). Column 3 (H1b) is a binary indicator equal to one in firm-months in which external write-access holders hold the majority of project gatekeeping; it uses the 386-firm subsample for which gatekeeping decisions are observable in the GitHub event archive. Columns 4 and 5 are log-of-one-plus transformations of monthly firm-level counts of external forks of the firm's repositories and external pull requests received. Controls comprise three time-varying firm-level indicators: venture-capital backing, public listing, and prior acquisition. Standard errors clustered at the firm level appear in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4. Mediation analysis. Governance change as the channel from shock to community engagement.

	Permissive project release (H1a)		Gatekeeping openness (H1b)	
Dependent variable:	Forks	Ext. PRs	Forks	Ext. PRs
Panel A. Baron-Kenny mediation paths				
Path a (Treatment → Mediator)	0.0709^{***}		0.1568^{***}	
	(0.0250)		(0.0554)	
Path b (Mediator → Dep. var.)	0.3234^{***}	0.5649^{***}	0.2105^{***}	0.3564^{***}
	(0.0396)	(0.0583)	(0.0585)	(0.0830)
Indirect effect ($a \times b$)	0.0229	0.0401	0.0330	0.0559
Sobel z	2.68	2.72	2.22	2.36
Sobel p-value	0.007	0.007	0.026	0.018
Share of total effect mediated	21.7%	39.6%	10.3%	24.3%
Panel B. Bias-corrected bootstrap 95% CI (5,000 cluster replications)				
Bootstrap point estimate	0.0229	0.0401	0.0330	0.0559
Bias-corrected 95% CI	[0.008, 0.043]	[0.014, 0.073]	[0.010, 0.072]	[0.017, 0.120]
CI excludes zero?	Yes	Yes	Yes	Yes
Panel C. Sample				
Observations	36,337	36,337	11,923	11,923
Firms (clusters)	664	664	386	386

Notes. This table reports mediation analysis for two governance mediators of the export-control shock's effect on community engagement: permissive project release (H1a), the log of one plus the rolling six-month count of newly released repositories under permissive licenses; and gatekeeping openness (H1b), a binary indicator equal to one in firm-months in which external developers hold the majority of active write-access positions. Each mediator is tested against each of the two community-engagement dependent variables, forks and external pull requests, yielding four single-mediator analyses. Panel A follows the three-equation framework of Baron and Kenny (1986). Path a regresses the mediator on Treated × Post; its estimate is shared across the two dependent-variable columns for each mediator. Path b regresses the dependent variable on the mediator, controlling for Treated × Post. The indirect effect is the product of a and b. The Sobel z-statistic and p-value test the indirect effect against the null of zero under normality. Because the product distribution is right-skewed, Panel B reports bias-corrected 95% confidence intervals from a cluster-resampled bootstrap of 5,000 replications (Preacher and Hayes, 2008); clusters are firms. Panel C gives the estimating samples: the permissive-release analyses use the full 664-firm matched sample, and the gatekeeping-openness analyses use the 386-firm gatekeeping-observable subsample. All regressions include firm and calendar-month fixed effects; controls comprise venture-capital backing, public listing, and prior acquisition. Standard errors clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. Robustness. Alternative matching specifications across the four hypothesis dependent variables.

	(1) Baseline CEM	(2) CEM alt. 1	(3) CEM alt. 2	(4) CEM alt. 3
Permissive project release (H1a)				
Treated × Post	0.0709***	0.0640***	0.0642***	0.0571**
	(0.0250)	(0.0240)	(0.0242)	(0.0234)
Gatekeeping openness (H1b)				
Treated × Post	0.1568***	0.1560***	0.1669***	0.1603***
	(0.0554)	(0.0561)	(0.0562)	(0.0554)
Forks (H2a)				
Treated × Post	0.1056**	0.0965**	0.0924**	0.0925**
	(0.0461)	(0.0446)	(0.0452)	(0.0440)
External pull requests (H2b)				
Treated × Post	0.1013**	0.0971**	0.0985**	0.0934**
	(0.0446)	(0.0421)	(0.0424)	(0.0404)
Firm fixed effects	Yes	Yes	Yes	Yes
Calendar-month FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Firms (H1a, H2a, H2b)	664	743	727	826
Firms (H1b subsample)	386	423	416	459
Matching variables	Age_{Q4}, GHAge_{Q4}, VC, Public, Size_{sm}	Age_{Q4}, GHAge_{Q4}, VC, Public	Age_{Q3}, GHAge_{Q3}, VC, Public, $\text{Size}_{\text{3-lev}}$	Age_{Q3}, GHAge_{Q3}, VC

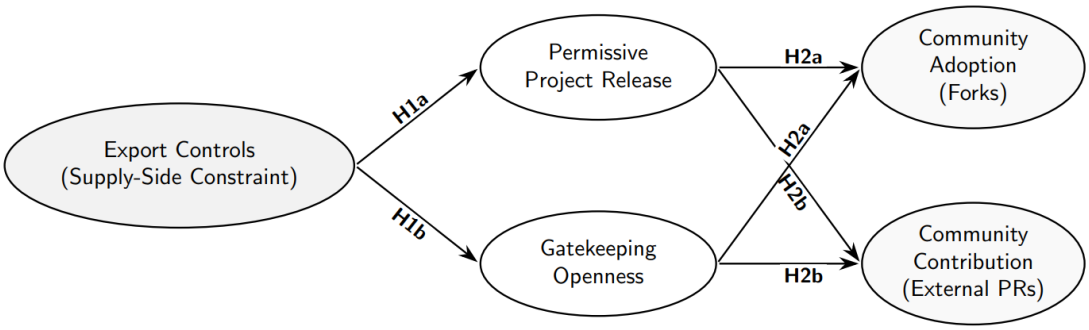
Notes. Each cell reports the Treated × Post coefficient from a separate difference-in-differences regression with firm and calendar-month fixed effects and the three time-varying firm-level controls (venture-capital backing, public listing, prior acquisition). Standard errors clustered at the firm level appear in parentheses. Column 1 reproduces the Baseline CEM used in Table 3: quartile bins on firm age and GitHub organization age, and exact matching on venture-capital backing, public listing, and the medium-or-large firm indicator (≥ 51 employees). Columns 2–4 vary the Coarsened Exact Matching (Iacus, King, & Porro, 2012) coarsening scheme. CEM alternative 1 drops the size exact-match but retains quartile age bins. CEM alternative 2 switches to tercile bins and replaces the binary size indicator with a three-level size variable. CEM alternative 3 uses tercile bins and matches exactly on venture-capital backing only. The firms-row counts differ across specifications because the coarsening grid changes which firms can be matched. The gatekeeping row reports the firm count for the H1b subsample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Cross-country evidence: India placebo and triple-difference specification

	(1) Perm. release (H1a)	(2) Gatekeep. (H1b)	(3) Forks (H2a)	(4) Ext. PR (H2b)
Panel A: India placebo				
Treated × Post	0.0010	0.0357	−0.0072	−0.0034
	(0.0047)	(0.0438)	(0.0062)	(0.0089)
Observations	360,102	41,886	360,102	360,102
Firms	5,994	1,610	5,994	5,994
Panel B: Chinese + Indian triple-difference				
Treated × Post × China	0.0659**	0.1280*	0.1160**	0.0912**
	(0.0255)	(0.0707)	(0.0462)	(0.0448)
Firm FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	396,439	53,809	396,439	396,439
Firms	6,658	1,996	6,658	6,658

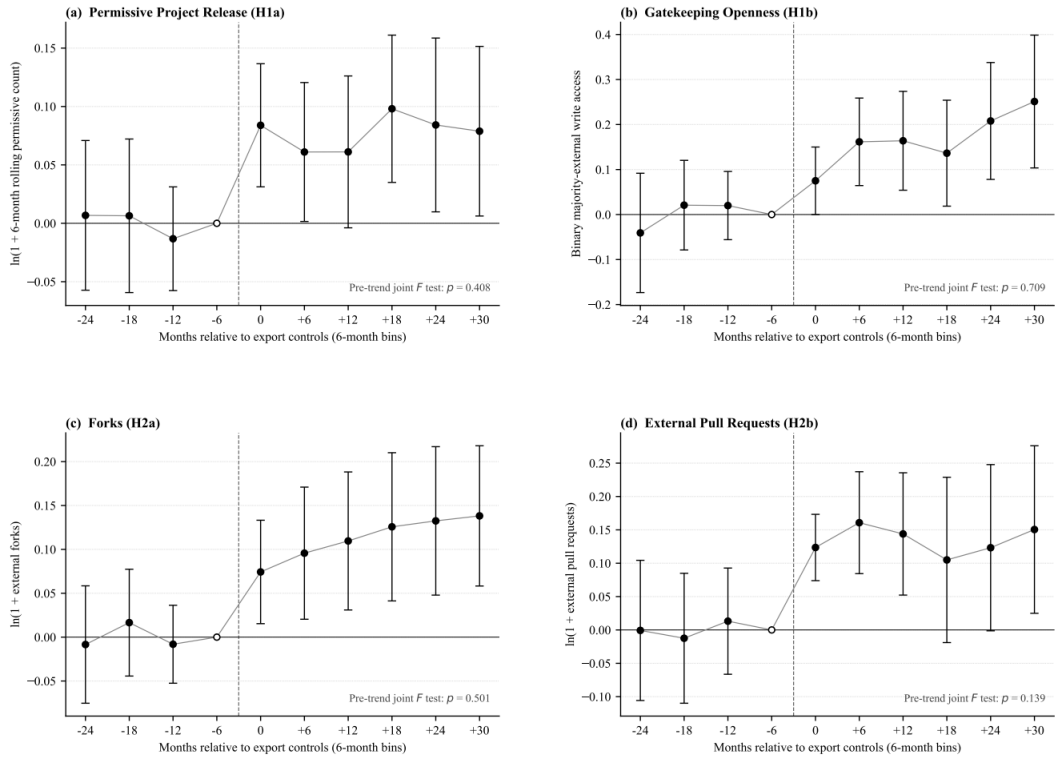
Notes. Panel A reports coefficients on Treated × Post from re-estimating the four hypothesis-test specifications on a comparable panel of Indian software firms. India was not subject to the October 2022 export controls, so any non-zero coefficient would suggest global trends rather than policy effects. All four coefficients are statistically indistinguishable from zero. Panel B reports the coefficient on Treated × Post × China from the triple-difference specification pooling Chinese and Indian matched panels. This triple interaction is the coefficient of interest: Chinese AI firms increased open-source activity relative to the combined comparison group of Indian AI firms, Chinese non-AI firms, and Indian non-AI firms. Standard errors clustered at the firm level are in parentheses. Column 2 (gatekeeping openness) restricts to firms with gatekeeping data. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1. Conceptual model of dependence reconfiguration through open-source governance.



Notes. The figure summarises the study's hypotheses as a directed graph. The October 2022 U.S. export controls on advanced AI computing hardware are the supply-side shock (left). H1a and H1b predict that the governance of firms' new open-source projects shifts toward a higher rate of permissive project release and toward external majorities in gatekeeping roles, respectively. H2a and H2b predict that the two governance shifts mediate the effect of the shock on community adoption (forks) and community contribution (external pull requests), respectively. Solid arrows denote causal paths tested in the main specification.

Figure 2. Event-study coefficients around the October 2022 export controls.



Notes. Each panel plots coefficients on Treated × Post by event-time bin around the October 2022 U.S. export controls on advanced AI computing hardware. The specification follows Borusyak, Jaravel, and Spiess (2024). Event time is measured in six-month bins; the bin $[m = -6, -1]$ is the omitted reference and its coefficient is pinned to zero by construction (hollow circle). Whiskers show 90 percent confidence intervals based on standard errors clustered at the firm level. The dashed vertical line marks the policy announcement. Joint F-tests of the pre-shock bin coefficients yield p-values of 0.408 (permissive project release), 0.709 (gatekeeping openness), 0.501 (forks), and 0.139 (external pull requests), all above the 0.10 threshold, consistent with the parallel-trajectories assumption underlying the main specification. All regressions include firm and calendar-month fixed effects and three time-varying firm-level controls: venture-capital backing, public listing, and prior acquisition. Panels (a), (c), and (d) use the 664-firm Coarsened Exact Matching sample (36,337 firm-month observations); panel (b) uses the 386-firm gatekeeping-observable subsample (11,923 firm-month observations).

Appendix A: Classification of Firms by the Level of AI Involvement

Identifying the treated and control groups for our analysis requires a reproducible procedure that separates software firms whose core products depend on advanced AI computing hardware from otherwise comparable software firms that do not. We use a large language model to assign each firm a level on a five-point AI value-chain rubric, following recent work that validates LLM-based text annotation against human coders (Gilardi et al. 2023). The classifier reads three inputs for each firm: the firm's name, its full Crunchbase business description, and its Crunchbase industry categories. It returns one of five levels together with a short textual justification. To address the concern that firms may have rebranded around AI after the October 2022 export controls, we run the classifier on two versions of each firm's Crunchbase profile: a 2022 snapshot and a 2025 snapshot. Treatment is assigned from the 2022 label, because the treatment indicator must reflect what the firm was at the time the policy took effect.

The rubric assigns each firm to one of the following five levels:

- *Non-AI* firms (level 1) have no meaningful AI component in their core product or operations.
- *AI feature* firms (level 2) add an AI capability to a primary product that is not itself centered on AI; an e-commerce platform that ships a chatbot customer-service feature is one example.
- *AI application* firms (level 3) offer a primary product whose value proposition is an AI capability delivered to end users.
- *AI infrastructure* firms (level 4) build the tools that support AI development by other firms, including model-serving platforms, vector and training-data pipelines, and developer tooling for AI workloads.
- *Foundation-model developer* firms (level 5) train and release large general-purpose models, such as large language models and diffusion models, as their primary product.

Firms at levels 3 through 5 form the treated group, because their core products depend directly on the advanced computing hardware restricted by the October 2022 controls. Firms at levels 1 and 2 form the control group.

The classifier is Qwen3 8B, an 8-billion-parameter open-weights large language model. It is served locally and queried at low temperature (0.2) for near-deterministic output, with up to two retries if the model returns output that fails schema validation. Each firm is presented in a single-shot prompt that states the task, lists the five levels with definitions, and provides level-specific text cues. For example, words such as "train," "pretrain," "foundation model," and "LLM" are listed as indicators of level 5; "MLOps," "vector database," and "RAG" are listed as indicators of level 4. The model returns a structured object containing the level, a primary-role tag, a compute-exposure tag, two to five evidence phrases quoted from the firm's own text, and a confidence score. Two firms illustrate the endpoints of the rubric. DeepSeek is categorized by Crunchbase under Artificial Intelligence, Generative AI, and Machine Learning, with a description that begins "DeepSeek is an artificial intelligence software company that focuses on developing open-source large language models (LLMs)." It receives level 5 with confidence 95, with evidence phrases on open-source LLM releases and with a compute-exposure tag of "high." Zhihu is categorized under social media and internet services, with a description that begins "Zhihu is a Chinese social network platform dedicated to questioning and answering." It receives level 1 with confidence 100,

with evidence phrases on the question-and-answer format and with a compute-exposure tag of "low."

Running the classifier on the 2022 and 2025 snapshots gives two labels per firm. We manually inspected the LLM's output against the underlying Crunchbase text and adjusted the assigned level where appropriate. Of the 1,974 firms, 1,684 labels were unchanged, 287 were upgraded, and three were downgraded; only seven of these adjustments crossed the binary treatment threshold (six upgrades plus one downgrade), so the LLM and manual review agreed on the treatment assignment for 99.6 percent of the sample. Smaller within-class adjustments are captured in the diff between the pre-review and post-review label files.

Five firms (Closeli, MetaX, DP Technology, Data Grand, and Arcvideo Tech) receive a level 1 or level 2 label in both snapshots but are nonetheless excluded from the analytic sample. Their 2022 Crunchbase descriptions are silent on AI: the texts describe the firms, respectively, as a generic IT company, a molecular simulation technology firm, an integrated circuit design firm, a semantic analysis developer, and an intelligent video solutions firm with no AI categories. Their 2025 Crunchbase descriptions document substantive AI activity: computer vision products, an AI-for-science foundation model, GPUs for AI training, natural language processing with robotic process automation, and face recognition with AI, respectively. Because we cannot confirm that any of these firms was a genuine control firm at the time of the shock, assigning them to the treated group would misrepresent the group they belonged to in October 2022, and retaining them as controls would allow post-shock rebranding to contaminate the counterfactual. We therefore drop all five rather than assign them to either side of the binary indicator.

After these steps, and after further dropping 17 firms founded after October 2022 that could not have responded to the policy, the final estimating sample comprises 143 treated firms at the AI application level or higher (levels 3 through 5) and 521 control firms at the AI feature or non-AI levels (levels 1 and 2).