

Chipmakers, Cloud Providers, and Frontier AI Labs: Circular Investment and the Economics of the Compute Race

Nianzu Xiong

Department of Economics, University of Oxford

October 2026

Abstract

Chipmakers and cloud providers take equity stakes in the frontier AI labs that buy their compute, and labs receive equity or warrants in their suppliers in return. I model this circular investment in an AI stack where labs race for the frontier in an all-pay contest with knowledge spillovers α and buy compute from suppliers that earn a margin. I obtain closed-form equilibria. A supplier's stake in a lab is neutral for compute; the lab's stake in its supplier acts as a volume rebate, which the supplier's stake amplifies. The value of the ownership network moves with compute at rate $\alpha + \mu_N - 1$, where μ_N is the share of each compute dollar retained as profit inside the network. Below this threshold, an impossibility theorem holds: no circular arrangement raises compute and benefits its participants. Even then, insiders can gain at the expense of the labs' outside investors. Exclusive deals discourage rival labs and create value exactly when inclusive deals destroy it. Circularity amplifies exposure to margin compression. With data-center and power layers, the thresholds acquire cross-layer feedback terms, and bringing power into the network can reverse the regime.

Keywords: circular investment, vendor financing, cross-shareholding, common ownership, AI, data centers, compute, all-pay contests, spillovers.

JEL: D43, D44, G32, L13, L22, L63, O31.

Contents

1	Introduction	4
2	Circular investment in the AI stack	6
3	The model	7
3.1	The stack	7
3.2	Ownership	9
4	A tool: all-pay contests with linear payoffs	10
5	Equilibrium: the role of reciprocal stakes	11
6	Value: a single threshold	13
6.1	The network-value identity	13
6.2	Instrument irrelevance	14
7	Who bears the losses? Outside investors	15
8	Exclusive versus broad-based investment	15
9	Margins and fragility	19
10	Designing circular deals	20
11	Many labs	21
12	Financing constraints: when one-way stakes matter	21
13	Data centers and power: a two-layer stack	22
13.1	Setup	22
13.2	Equilibrium and value	23
13.3	Power and leakage	24
14	Welfare and policy	25
14.1	Private versus social thresholds	25
14.2	Policy implications	25

15 An illustrative calibration	26
16 Discussion of modeling choices	27
17 Testable implications	27
18 Conclusion	28
A Proofs	29
B Numerical verification	29

1 Introduction

The frontier of artificial intelligence is being built on an unusual web of financial ties. Chip designers and cloud providers invest in the AI labs that buy their products. The labs commit to multi-year, multi-gigawatt compute purchases, and in some deals they receive equity or warrants in their suppliers that vest as purchases are made. Investors and journalists call these arrangements “circular”: money flows from supplier to lab as equity, back from lab to supplier as revenue, and sometimes back again as equity in the supplier ([The Register, 2025](#); [Morningstar, 2025](#)). Two questions follow. Do these arrangements raise investment in compute? And do they create value, or do they inflate compute spending at someone else’s expense?

This paper builds a model to answer both. The approach extends the framework of [Xiong \(2026\)](#), which studies R&D contests with spillovers and symmetric cross-shareholdings across industries and proves that no symmetric cross-shareholding structure can raise both R&D and investors’ profits. That result rests on two facts: winner-take-all contests dissipate rents, and an investor whose portfolio weights sum to one earns the profit of a representative firm. Circular investment in the AI stack differs in two ways that matter. The ties are *vertical*: suppliers earn a margin on every dollar labs spend. And they are *asymmetric and partial*: a supplier may back one lab or all of them, and labs are largely funded by outside investors. Both features turn out to be decisive.

The model in brief. Two AI labs race for the frontier. Each spends on compute; the lab that spends more wins a prize V that grows with total compute spending at rate α , capturing knowledge diffusion and market expansion. A dollar of lab spending earns a margin μ for the compute layer (chips and cloud). Ownership links the layers. The compute supplier’s controlling owner holds a share σ of each lab. Each lab’s insiders hold a share ω of the supplier, which captures equity or warrants that labs receive. Lab insiders may also hold a share h of the rival lab, and outside investors hold shares ϕ_L of each lab and ϕ_S of the supplier. In an extension, a second layer of data-center and power developers also races to build capacity, and the two layers’ investments spill over onto each other.

Results. All equilibria are in closed form. The tool is a lemma for symmetric all-pay contests in which payoffs are linear in both players’ efforts ([Lemma 1](#)); it covers any ownership structure in the model.

1. *Reciprocal stakes matter* ([Theorem 1](#), [Proposition 1](#)). A supplier’s stake in a lab, on its own, has no effect on compute. Lab insiders still bear the full marginal cost of compute and still receive the full marginal benefit, scaled by the same share. What raises compute is the lab’s stake in the supplier. Lab insiders get back $\omega\mu$ of every dollar spent, a volume rebate of $r = \omega\mu/\theta$ per dollar of their own cost, where θ is their stake in the lab. The supplier’s stake σ and outside financing ϕ_L shrink θ and thereby amplify the rebate. Equilibrium compute is $\bar{e} = \bar{e}^0(1 - \alpha)/(1 - \alpha - r)$, a circularity multiplier. Labs post accounting losses while their insiders gain through their stakes in

the supplier.

2. *A single value threshold (Theorem 2).* Margins are transfers within the network. By accounting, the combined value of all firms linked by ownership is $J = V - (1 - \alpha - \mu_N) T$, where T is total compute and μ_N is the in-network margin per dollar. Every arrangement, whether equity, warrants, purchase commitments or rebates, changes network value by $(\alpha + \mu_N - 1)\Delta T$. If $\alpha + \mu_N < 1$, a *circular-investment impossibility theorem* holds: no arrangement that raises compute can make its participants better off, even with side payments. If $\alpha + \mu_N > 1$, the labs under-invest from the network's point of view because they pay margins that are only transfers, and every compute-raising arrangement creates value.

3. *The incidence of losses (Theorem 3).* With outside investors, the insiders' gain from extra compute is proportional to $\mu(1 - \phi_S) - (1 - \alpha)(1 - \phi_L)$. When labs are mostly financed from outside, insiders can profit from expansions of compute that reduce total value. The losses are borne by the labs' outside shareholders, through the labs' accounting losses.

4. *Exclusive versus broad-based investment (Theorem 4).* An exclusive circular deal with one lab creates an asymmetric contest, which I solve in closed form with endogenous prizes. Notably, the favored lab's expected spending does not change: the deal works only by discouraging the rival, whose spending falls and who gives up the frontier more often. Total compute falls. Exclusive deals therefore create value exactly when $\alpha + \mu_N < 1$, the region in which inclusive deals destroy it. Suppliers with high margins and large spillovers should invest across all labs; investors with low margins, and vertically integrated firms, should concentrate on one lab.

5. *Fragility (Proposition 5).* Circularity raises the elasticity of compute spending with respect to the supplier's margin, from zero to $r/(1 - \alpha - r)$, and makes the supplier's revenue more than proportionally sensitive. Structures built when margins exceed $1 - \alpha$ become value-destroying if competition compresses margins below $1 - \alpha$, and the more circular the structure, the larger the loss.

6. *Data centers and power (Theorem 5, Proposition 9).* When a second layer of data-center and power developers also races, and the layers' investments spill over onto each other, labs' stakes in infrastructure create value iff $\alpha + \mu + \beta_L \beta_P / (1 - \alpha_P) > 1$. Infrastructure firms' stakes in labs create value iff $\alpha_P + \beta_L(\mu + \beta_P) / (1 - \alpha) > 1$. The β terms are cross-layer feedback loops. Margins earned outside the network, such as those of independent power producers, are invisible to it; bringing power into the network raises μ_N and can flip the regime.

Welfare and policy (Section 14). A planner counts all margins as transfers and adds the external value of compute η (positive for diffusion to users, negative for grid congestion, emissions or safety risk). Privately profitable inclusive deals are then socially beneficial whenever $\eta \geq 0$. Exclusive deals can be privately profitable but socially harmful: they lower compute and concentrate the frontier.

Related literature. The contest machinery builds on complete-information all-pay auctions (Hillman and Riley, 1989; Baye et al., 1996; Siegel, 2009) and on contests whose prizes depend on effort (Baye et al., 2012). The asymmetric equilibrium with an endogenous common prize in Section 8 appears to be new in closed form. The ownership side relates to the literatures on partial cross-ownership (Reynolds and Snapp, 1986; Gilo et al., 2006), partial vertical ownership and foreclosure (Flath, 1989; Greenlee and Raskovich, 2006; Levy et al., 2018; Ordover et al., 1990), vendor financing (Brennan et al., 1988), and common ownership and innovation (López and Vives, 2019; Antón et al., 2021). The value thresholds generalize the comparison between noncooperative and cooperative R&D in d’Aspremont and Jacquemin (1988) and Kamien et al. (1992) to vertical margins and cross-layer spillovers. Relative to Xiong (2026), the contributions are the vertical margin channel, asymmetric (exclusive) holdings, outside investors, and the multi-layer stack.

Roadmap. Section 2 summarizes the arrangements that motivate the model. Section 3 presents the model and Section 4 the contest lemma. Sections 5–9 contain the main results for a single stack. Section 13 adds data centers and power. Sections 14–15 discuss welfare, policy and an illustrative calibration. Proofs not given in the text are in Appendix A; numerical verification is in Appendix B.

2 Circular investment in the AI stack

This section describes the arrangements that motivate the model. Terms of individual deals change frequently and are often only partly disclosed. The model is built to capture the *types* of linkage rather than any particular contract.

Supplier equity in labs. Chip designers and cloud providers hold equity in AI labs. Microsoft has been OpenAI’s largest outside investor since 2019 and, after OpenAI’s 2025 recapitalization, holds a stake of roughly 27% (Business Standard, 2025). In November 2025, Microsoft and Nvidia announced investments in Anthropic of up to \$15 billion combined, alongside an Anthropic commitment to purchase \$30 billion of Azure compute (Pure AI, 2025). Nvidia’s equity investments in AI companies in 2026 have reportedly exceeded \$40 billion, anchored by a \$30 billion investment in OpenAI (The Next Web, 2026). In the model, these are the supplier’s stakes σ .

Lab equity in suppliers. In October 2025, AMD and OpenAI announced a partnership to deploy six gigawatts of AMD GPUs; OpenAI received a warrant for AMD shares that vests as deployment milestones are met (AMD, 2025). A warrant that vests with purchases is, economically, a volume rebate paid in the supplier’s equity: the more the lab buys, the more of the supplier’s margin it captures. In the model, these are the lab’s stakes ω .

Compute and data-center layers. Data-center operators sit between chip makers and labs. Nvidia has invested repeatedly in CoreWeave, including a further \$2 billion in January 2026 (Bloomberg, 2026). The result is a chain of stakes and purchase commitments running from chips through data centers to labs and back (The Register, 2025).

Power. The binding constraint on new data centers is increasingly electricity. Hyperscalers have signed long-term power purchase agreements and made commitments to nuclear and other firm generation. Most of these are contracts rather than equity, and their realism has been questioned (Carnegie, 2026). In the model, power is a layer whose margins are either inside or outside the ownership network.

The debate. Commentators disagree on whether circularity is benign (risk sharing, financing of a capital-intensive buildout) or a source of fragility (revenue round-tripping, inflated demand) (Morningstar, 2025; The Register, 2025). The model gives conditions under which each view is right. Figure 1 summarizes the links described above.

3 The model

3.1 The stack


There are two AI labs, L_1 and L_2 ; a compute supplier S , representing the chips-and-cloud layer; and an outside layer O , representing, for example, independent power producers. Section 13 adds a second racing layer.

The frontier race. Lab i chooses compute spending $e_i \geq 0$, measured in dollars. The lab that spends more wins the frontier and earns

$$R(e_1, e_2) = V + \alpha(e_1 + e_2), \quad V > 0, \alpha \in [0, 1]. \quad (1)$$

Ties are broken by a fair coin. The loser earns nothing from the race. V is the value of leadership: the rents of the frontier model over its product cycle. The spillover α captures the fact that the winner's prize grows with both labs' compute: some of the loser's work diffuses (through publications, staff mobility, distillation or open weights), and both labs' spending expands the market for frontier AI. The race is the all-pay contest of Baye et al. (1996) with the endogenous prize of Baye et al. (2012), as in Xiong (2026).

Margins. A dollar of lab spending generates profit $\mu \in [0, 1)$ for the supplier S and profit $\mu_O \geq 0$ for the outside layer O . The rest, $1 - \mu - \mu_O > 0$, is real resource cost. So $\Pi_S = \mu(e_1 + e_2)$ and



—————> equity (holder → firm) > warrant (holder → issuer) - - -> compute purchases (buyer → seller)

- | | |
|---|---|
| 1 Nvidia's \$30bn investment in OpenAI (2026) | 6 OpenAI's AMD warrant, vesting with deployment |
| 2 Nvidia's investment in Anthropic (Nov. 2025) | 7 OpenAI's 6 GW AMD GPU deployment |
| 3 Microsoft's $\approx 27\%$ stake in OpenAI | 8 OpenAI's purchases of Azure compute |
| 4 Microsoft's investment in Anthropic (Nov. 2025) | 9 Anthropic's \$30bn Azure commitment |
| 5 Nvidia's investments in CoreWeave | 10 CoreWeave's purchases of Nvidia GPUs |

Figure 1: A stylized map of selected circular links in the AI stack, based on the public reports cited in Section 2 (AMD, 2025; Business Standard, 2025; Pure AI, 2025; The Next Web, 2026; Bloomberg, 2026). Circles are firms: blue for chipmakers, orange for cloud and data-center providers, green for frontier labs. The map is illustrative and not exhaustive; arrows are not to scale, and many terms are only partly disclosed. The model maps solid arrows from suppliers into labs to σ , dotted arrows from labs into suppliers to ω , and dashed arrows to the spending e_i on which suppliers earn the margin μ .

$\Pi_O = \mu_O(e_1 + e_2)$, and lab i 's profit is

$$\Pi_i = R \cdot \mathbf{1}\{i \text{ wins}\} - e_i.$$

Margins are taken as given; Section 9 studies changes in them.

3.2 Ownership

Each firm has a controlling *insider*: the lab insiders 1,2 and the supplier insider S . Insiders and passive outside investors hold the following shares:

	Lab i	Lab j	Supplier S
Lab- i insider	θ	h	ω
Supplier insider	σ	σ	$1 - 2\omega - \phi_S$
Outside investors	ϕ_L	ϕ_L	ϕ_S

Each column sums to one, so $\theta = 1 - \sigma - h - \phi_L$. Insiders control their firms' compute decisions (for instance through dual-class shares or founder control) even when their economic stake is below one half. I call (σ, ω) the *circular* holdings, h the *horizontal* holding, and (ϕ_L, ϕ_S) *outside finance*. The *network* is the set of insiders. Its in-network margin is $\mu_N = \mu$, since S is inside and O outside. Section 13 brings additional layers in. Figure 2 depicts the structure.

Decision rights. Each lab's spending is chosen to maximize its insider's expected payoff:

$$W_i = \theta \Pi_i + h \Pi_j + \omega \Pi_S. \quad (2)$$

The supplier insider's payoff is $W_S = (1 - 2\omega - \phi_S)\Pi_S + \sigma(\Pi_1 + \Pi_2)$. Product-market conduct after the race is summarized by V and does not depend on ownership. This isolates the investment channel; Section 14 discusses conduct effects.

Instruments. The share ω stands for any arrangement through which lab insiders capture part of the supplier's margin: equity, warrants vesting with purchases, or revenue-share rebates. Section 6 shows that purchase commitments and other instruments are evaluated by the same value threshold.

Timing. Holdings are in place. Labs choose spending simultaneously. The winner is determined and payoffs are realized. The benchmark has no holdings: $\sigma = \omega = h = \phi_L = \phi_S = 0$ and $\theta = 1$.




Figure 2: Ownership and cash flows in the model. The supplier insider holds σ of each lab; each lab insider holds ω of the supplier and h of the rival lab; outside investors hold ϕ_L of each lab and ϕ_S of the supplier (only one ϕ_L arrow is drawn). Each dollar of lab spending yields margin μ to the supplier and μ_O to the outside power layer. The ownership network consists of the three insiders; O is outside it unless linked by equity (Section 13).

4 A tool: all-pay contests with linear payoffs

All results rest on the following lemma, which generalizes the standard symmetric all-pay equilibrium to payoffs linear in both players' efforts.

Lemma 1 (Linear contests). *Two players choose efforts $e_i, e_j \geq 0$. If $e_i > e_j$, player i 's payoff is $w_0 + w_1 e_i + w_2 e_j$; if $e_i < e_j$, it is $\ell_0 + \ell_1 e_i + \ell_2 e_j$; ties are split. Suppose*

$$\Delta \equiv w_0 - \ell_0 > 0, \quad d \equiv w_1 - \ell_1 = w_2 - \ell_2 > 0, \quad \kappa \equiv -\ell_1 > d.$$

Then there is a symmetric equilibrium in which each player draws effort from

$$F(e) = \frac{\kappa}{d} \left[1 - \left(\frac{\Delta}{\Delta + 2de} \right)^{1/2} \right], \quad e \in [0, b], \quad b = \frac{\Delta}{2d} \left[\left(\frac{\kappa}{\kappa - d} \right)^2 - 1 \right]. \quad (3)$$

Expected effort and the expected equilibrium payoff are

$$m = \frac{\Delta}{2(\kappa - d)}, \quad U^* = \ell_0 + \ell_2 m. \quad (4)$$

When $d = 0$, F is uniform on $[0, \Delta/\kappa]$ and (4) still holds.

Proof. Let the opponent draw from an atomless c.d.f. F on $[0, b]$ with density f . The expected payoff

from $e \in [0, b]$ is

$$U(e) = \int_0^e (w_0 + w_1 e + w_2 y) dF(y) + \int_e^b (\ell_0 + \ell_1 e + \ell_2 y) dF(y).$$

Differentiating and using the hypotheses,

$$U'(e) = f(e) [\Delta + 2de] + dF(e) - \kappa.$$

Setting $U' = 0$ and substituting $u = \Delta + 2de$ gives $2uF'(u) + F = \kappa/d$, i.e. $\frac{d}{du}(F\sqrt{u}) = \frac{\kappa}{2d\sqrt{u}}$. With $F = 0$ at $u = \Delta$, this yields (3). F is strictly increasing and reaches 1 at b because $\kappa > d$. So U is constant on $[0, b]$. For $e > b$ the player wins surely and $U(e) = w_0 + w_1 e + w_2 m$, which is strictly decreasing because $w_1 = \ell_1 + d = -(\kappa - d) < 0$. Hence F is a best response to itself.

For the mean, integrate $f u + dF = \kappa$ over $[0, b]$. Since $\int_0^b f u de = u(b) - 2d \int_0^b F$, this gives $u(b) - d \int_0^b F = \kappa b$. Then $m = b - \int_0^b F = [(\kappa - d)b - \Delta]/d$. Substituting b and using $\kappa^2 - (\kappa - d)^2 = d(2\kappa - d)$ gives $m = \Delta/[2(\kappa - d)]$. Finally, $U^* = U(0)$: at $e = 0$ the player loses with probability one and earns $\ell_0 + \ell_2 m$. $\square$

The lemma has a simple economic reading. $\Delta + d(e_i + e_j)$ is the *value of winning over losing*, and $\kappa - d$ is the *net marginal cost of effort*: the marginal cost when losing, less the part recovered through the prize. Expected effort is half the intercept of the value of winning, divided by the net marginal cost. Payoffs come from what a player earns when it loses with certainty, because the rents from winning are competed away.

5 Equilibrium: the role of reciprocal stakes

Theorem 1 (Equilibrium with circular holdings). *Suppose $\alpha > 0$, $\theta > h$ and $\theta(1 - \alpha) > \omega\mu$. The lab race has a symmetric equilibrium of the form (3), with*

$$\Delta = (\theta - h)V, \quad d = (\theta - h)\alpha, \quad \kappa = \theta - h\alpha - \omega\mu.$$

Expected spending per lab is

$$\bar{e} = \frac{(\theta - h)V}{2[\theta(1 - \alpha) - \omega\mu]} = \frac{(1 - h/\theta)V}{2[1 - \alpha - r]}, \quad r \equiv \frac{\omega\mu}{\theta}, \quad (5)$$

and each lab insider's expected payoff is

$$\mathbb{E}W_i = hV + [\omega\mu - h(1 - \alpha)] \bar{e}. \quad (6)$$

In the benchmark, $\bar{e}^0 = V/[2(1 - \alpha)]$; labs earn zero, and the supplier earns $\mu V/(1 - \alpha)$.

Proof. By (2), if lab i wins, insider i earns $\theta(R - e_i) - he_j + \omega\mu(e_i + e_j)$; if it loses, $-\theta e_i + h(R -$




Figure 3: The circularity multiplier (Proposition 1), with $\alpha = 0.2$, $\mu = 0.6$ and $h = \phi_L = 0$. With $\omega = 0$, the supplier's stake σ has no effect. With $\omega > 0$, σ amplifies the effect of ω .

$e_j) + \omega\mu(e_i + e_j)$. Collecting terms gives $w_0 = \theta V$, $w_1 = \theta\alpha - \theta + \omega\mu$, $w_2 = \theta\alpha - h + \omega\mu$, $\ell_0 = hV$, $\ell_1 = h\alpha - \theta + \omega\mu$ and $\ell_2 = h\alpha - h + \omega\mu$. Hence $\Delta = (\theta - h)V$, $w_1 - \ell_1 = w_2 - \ell_2 = (\theta - h)\alpha$, and $\kappa = \theta - h\alpha - \omega\mu$, so $\kappa - d = \theta(1 - \alpha) - \omega\mu > 0$. Apply Lemma 1. $\square$

The parameter $r = \omega\mu/\theta$ is the *effective rebate*. For each dollar a lab spends, its insider bears θ dollars of cost (the rest falls on the other shareholders) and recovers $\omega\mu$ dollars through its stake in the supplier. Relative to its own cost, the insider gets back r per dollar.

Proposition 1 (Reciprocal stakes). *Let $h = 0$. Then:*

- (a) One-way stakes are neutral. If $\omega = 0$, $\bar{e} = \bar{e}^0$ for every σ and ϕ_L .
- (b) The reverse stake is a rebate. $\bar{e}$ is strictly increasing in ω and in μ whenever $\omega > 0$.
- (c) Amplification. If $\omega > 0$, $\bar{e}$ is strictly increasing in σ and in ϕ_L , because both shrink θ and raise r .
- (d) Multiplier. $\bar{e}/\bar{e}^0 = (1 - \alpha)/(1 - \alpha - r)$, which grows without bound as $r \uparrow 1 - \alpha$.

With $h > 0$, $\bar{e}$ is strictly decreasing in h .

Proof. Immediate from (5). $\square$

Part (a) is worth stressing. Equity alone does not buy compute. A supplier that takes a passive stake in a lab gives the lab cash, but it does not change the lab insiders' trade-off between the cost of compute and the chance of winning; both are scaled by the same θ . What changes behavior is a link that pays the lab *per unit of compute*: a stake in the supplier, a warrant that vests with purchases, a rebate. Once such a link exists, the supplier's stake in the lab (and the lab's reliance on outside capital) magnifies it, because the insiders' own share of the cost falls while their share of the supplier's margin does not. Circularity is multiplicative.

Proposition 2 (Accounting losses, insider gains). *Let $h = 0$ and $\omega > 0$. Each lab's expected accounting profit is*

$$\mathbb{E}\Pi_i = \frac{V}{2} - (1 - \alpha)\bar{e} = -\frac{rV}{2(1 - \alpha - r)} < 0,$$

while each lab insider's expected payoff is $\mathbb{E}W_i = \omega\mu\bar{e} > 0$. The supplier's expected profit is $2\mu\bar{e}$, which is increasing in r .

Proof. Taking expectations, $\mathbb{E}\Pi_i = \frac{1}{2}\mathbb{E}R - \bar{e} = \frac{V}{2} + \alpha\bar{e} - \bar{e}$; substitute (5). The insider payoff follows from (6). $\square$

Circular structures therefore produce a recognizable pattern: labs that lose money on paper, insiders who do well through their claims on the supplier, and a supplier whose revenue rises with the degree of circularity. Whether the system as a whole gains is a separate question, answered next.

Example 1. Let $V = 1$, $\alpha = 0.2$ and $\mu = 0.6$. In the benchmark, each lab spends $\bar{e}^0 = 0.625$, labs earn zero, and the supplier earns 0.75. Now let the supplier take $\sigma = 0.3$ of each lab and each lab's insiders receive $\omega = 0.2$ of the supplier (with $h = \phi_L = 0$). Then $\theta = 0.7$ and $r = 0.2 \cdot 0.6 / 0.7 = 0.171$, so spending rises to $\bar{e} = 0.795$, an increase of 27%. Each lab's expected accounting profit is -0.136 . Each lab insider earns $\omega\mu\bar{e} = 0.095$. The supplier's profit rises to 0.955, of which its insider retains 60%, plus its share σ of the labs' losses. Combined network value falls from 0.75 to 0.682, because $\alpha + \mu = 0.8 < 1$ (Theorem 2).

6 Value: a single threshold

6.1 The network-value identity

Consider the network of insiders, and suppose for now that there is no outside finance ($\phi_L = \phi_S = 0$). Ownership is then closed: every share of every network firm is held by an insider. The network's combined payoff is the sum of the firms' profits, $J = \Pi_1 + \Pi_2 + \Pi_S$.

Theorem 2 (Network value and the circular-investment impossibility theorem). *Let $T = \mathbb{E}(e_1 + e_2)$ be expected total compute. For any strategy profile and any arrangement among network members that leaves V , α and μ unchanged (equity stakes, warrants, rebates, purchase commitments, side payments),*

$$\mathbb{E}J = V - (1 - \alpha - \mu_N) T. \tag{7}$$

Hence, relative to the benchmark:

- (a) *Impossibility. If $\alpha + \mu_N < 1$, every arrangement that raises total compute lowers the network's value. No such arrangement can make all participants better off, even with side payments.*




Figure 4: Change in network value from a circular structure with $\sigma = 0.2$, as a function of ω ($\alpha = 0.2$). The sign is determined by $\alpha + \mu - 1$ (Theorem 2).

(b) Possibility. If $\alpha + \mu_N > 1$, every arrangement that raises total compute raises network value. With side payments, it can make all participants better off.

Proof. Exactly one lab wins, so $\Pi_1 + \Pi_2 = R - e_1 - e_2$ for every profile. Adding $\Pi_S = \mu_N(e_1 + e_2)$ and taking expectations gives (7). Side payments only redistribute J . So all participants can gain if and only if J rises, which requires $(\alpha + \mu_N - 1)\Delta T > 0$. $\square$

The identity is an accounting fact, but it organizes everything that follows. From the network's point of view, a dollar of compute costs $1 - \mu_N$ (the margin is a transfer to the supplier) and returns α in a larger prize. The labs, however, pay the full dollar. When $\alpha + \mu_N < 1$, the rent-dissipating race already over-invests from the network's point of view, and any circular arrangement that raises compute further makes things worse. This is the analogue, for vertical networks, of the impossibility theorem in Xiong (2026). When $\alpha + \mu_N > 1$, the margin separates the labs' private cost from the network's cost: labs under-invest because they pay a markup that is only a transfer within the network. Circular arrangements then correct a double-marginalization problem in compute, and the network gains.

6.2 Instrument irrelevance

Theorem 2 does not depend on how compute is raised. To illustrate, consider a take-or-pay purchase commitment, a common feature of supplier deals.

Proposition 3 (Purchase commitments). *Without holdings, suppose each lab commits to pay at least $C > 0$ for compute, with usage up to C costing nothing at the margin. Then each lab spends $\bar{e} = \bar{e}^0 + C/(1 - \alpha)$ in expectation, its expected accounting profit is $-C$, and network value changes by $-2(1 - \alpha - \mu_N)C/(1 - \alpha)$.*

Proof. Write $e_i = C + x_i$ with $x_i \geq 0$. The race over x has prize $V' + \alpha(x_1 + x_2)$ with $V' = V + 2\alpha C$ and unit marginal cost. By Lemma 1, $\mathbb{E}x = V'/[2(1 - \alpha)]$, so $\bar{e} = C + V'/[2(1 - \alpha)] = \bar{e}^0 + C/(1 - \alpha)$. The race over x dissipates its rents fully, leaving a profit of $-C$. Apply (7). $\square$

A commitment of C raises spending by more than C , because of spillovers: each lab's committed compute raises the prize for which both labs compete. Commitments, equity rebates and warrants are thus different instruments with the same value test. A commitment that comes bundled with an equity investment ("we invest, you commit to our cloud") creates network value if and only if $\alpha + \mu_N > 1$. The equity part only splits the surplus or the loss.

7 Who bears the losses? Outside investors

AI labs are financed largely by outside investors, and suppliers are publicly held. With outside finance, the insiders' combined payoff is $J_{\text{ins}} = J - \phi_L(\Pi_1 + \Pi_2) - \phi_S\Pi_S$.

Theorem 3 (Insiders versus outsiders). *Let $h = 0$. A circular structure that raises spending per lab by $\Delta\bar{e} > 0$ changes the insiders' combined payoff by*

$$\Delta J_{\text{ins}} = 2\Delta\bar{e} \left[\mu(1 - \phi_S) - (1 - \alpha)(1 - \phi_L) \right]. \quad (8)$$

The labs' outside investors lose $2\phi_L(1 - \alpha)\Delta\bar{e}$, and the supplier's outside investors gain $2\phi_S\mu\Delta\bar{e}$. In particular, if

$$\frac{(1 - \alpha)(1 - \phi_L)}{1 - \phi_S} < \mu < 1 - \alpha,$$

the structure is profitable for the insiders and destroys total value. The net loss is borne by the labs' outside investors.

Proof. By Theorem 2 and Proposition 2, $\Delta J = 2(\alpha + \mu - 1)\Delta\bar{e}$, $\Delta\mathbb{E}\Pi_i = -(1 - \alpha)\Delta\bar{e}$ and $\Delta\mathbb{E}\Pi_S = 2\mu\Delta\bar{e}$. Then $\Delta J_{\text{ins}} = \Delta J + 2\phi_L(1 - \alpha)\Delta\bar{e} - 2\phi_S\mu\Delta\bar{e}$, which rearranges to (8). $\square$

Theorem 3 explains why circular structures can be attractive to the people who design them even when they destroy value. The labs' accounting losses are shared with outside investors, while the margin they generate flows to the supplier and, through ω , back to the lab insiders. Outside finance enters twice. It raises the rebate $r = \omega\mu/\theta$ (Proposition 1(c)), and it shifts losses onto outsiders. The model thus identifies a precise channel through which an expansion of compute can be privately rational and collectively wasteful. It also identifies who bears the risk: the outside investors in the labs, not the suppliers' shareholders, who benefit from higher volume.

8 Exclusive versus broad-based investment

Some suppliers invest in several competing labs; others tie themselves to one. In the model, an *inclusive* deal gives both labs the same holdings; an *exclusive* deal links the supplier to lab 1 only. With $h = 0$, lab 1's insider then has rebate $r = \omega\mu/\theta$ and normalized marginal cost $c \equiv 1 - r \in (\alpha, 1)$,




Figure 5: Theorem 3 with $\alpha = 0.2$. Between the solid line and the dotted line ($\phi_S = 0$), compute-raising circular structures benefit insiders but destroy total value, at the expense of the labs' outside investors.

while lab 2 has marginal cost 1. The exclusive deal turns the race into an asymmetric contest with a common endogenous prize.

Lemma 2 (Asymmetric linear contest). *Two players compete for the prize $V + \alpha(e_1 + e_2)$, with $\alpha > 0$, at marginal costs $c_1 < c_2$, where $c_1 > \alpha$. There is an equilibrium in which both players randomize on $[0, b]$, with $b = \frac{V}{2\alpha} \left[\left(\frac{c_2}{c_2 - \alpha} \right)^2 - 1 \right]$ and $u \equiv V + 2\alpha e$:*

$$G_1(e) = \frac{c_2}{\alpha} \left[1 - \sqrt{V/u} \right], \quad G_2(e) = \frac{c_1}{\alpha} - \frac{(c_1 - \alpha)c_2}{\alpha(c_2 - \alpha)} \sqrt{V/u}.$$

The weak player has an atom at zero of mass $G_2(0) = \frac{c_2 - c_1}{c_2 - \alpha}$. Moreover:

- (i) expected efforts are $\mathbb{E}e_1 = \frac{V}{2(c_2 - \alpha)}$ and $\mathbb{E}e_2 = \frac{(c_1 - \alpha)V}{2(c_2 - \alpha)^2}$;
- (ii) the strong player wins with probability $\frac{2c_2 - c_1 - \alpha}{2(c_2 - \alpha)}$;
- (iii) equilibrium payoffs are $V \frac{c_2 - c_1}{c_2 - \alpha}$ for the strong player and 0 for the weak player.

Proof. See Appendix A. □

Two features of Lemma 2 drive the results. First, the strong player's expected effort depends only on the weak player's cost: in equilibrium, G_1 must keep the weak player indifferent. Second, the weak player responds to its disadvantage by staying out with positive probability.

Figure 6 illustrates. With $c_2 = 1$, the strong player's distribution G_1 coincides *exactly* with the symmetric benchmark distribution: a lab that receives an exclusive rebate behaves as if nothing had changed. The rival's distribution shifts down. It places mass $(1 - c_1)/(1 - \alpha)$ at zero and mixes over the same interval with lower density. All of the deal's effect on behavior thus operates through the rival.




Figure 6: Equilibrium spending distributions under an exclusive deal (Lemma 2), with $V = 1$, $\alpha = 0.2$, $c_1 = 0.75$ and $c_2 = 1$. The favored lab's distribution equals the symmetric benchmark. The rival stays out with probability 0.31, and the favored lab wins with probability 0.66.

Theorem 4 (Exclusive versus inclusive circularity). *Fix an effective rebate $r \in (0, 1 - \alpha)$ and let $c = 1 - r$.*

- (a) *Exclusive deal. The favored lab's expected spending is unchanged at $\bar{e}^0 = V / [2(1 - \alpha)]$. The rival's falls to $(c - \alpha)V / [2(1 - \alpha)^2]$, and the favored lab wins with probability $(2 - \alpha - c) / [2(1 - \alpha)] > 1/2$. Total compute falls by $(1 - c)V / [2(1 - \alpha)^2]$, and*

$$\Delta J^{\text{excl}} = (1 - \alpha - \mu_N) \frac{(1 - c)V}{2(1 - \alpha)^2}.$$

- (b) *Inclusive deal. Total compute rises by $(1 - c)V / [(c - \alpha)(1 - \alpha)]$, and*

$$\Delta J^{\text{incl}} = -(1 - \alpha - \mu_N) \frac{(1 - c)V}{(c - \alpha)(1 - \alpha)}.$$

- (c) *Hence $\text{sign } \Delta J^{\text{excl}} = \text{sign}(1 - \alpha - \mu_N) = -\text{sign } \Delta J^{\text{incl}}$. If $\alpha + \mu_N < 1$, a supplier and a lab gain from an exclusive deal and lose from investing in both labs. If $\alpha + \mu_N > 1$, the reverse holds.*

In both cases the unlinked lab earns zero, so ΔJ is also the coalition's gain.

Proof. Apply Lemma 2 with $c_1 = c$ and $c_2 = 1$ for (a), and Theorem 1 with r for (b). In both cases $J = V - (1 - \alpha - \mu_N)T$, by the argument of Theorem 2. Lab 2 earns zero in the benchmark, as a weak player in the exclusive case, and in the symmetric case (where it is a coalition member anyway). $\square$

The exclusive deal is a pure discouragement device. It does not make the favored lab spend more; it makes the rival spend less and concede more often. Total compute falls, which helps the network precisely when compute is over-supplied ($\alpha + \mu_N < 1$). Inclusive deals do the opposite.

The theorem therefore offers a rationale for two observed patterns: some high-margin component suppliers spread their investments across many labs, while cloud providers that also own or sponsor labs tend to concentrate on one. In the model, the choice is governed by a single number, $\alpha + \mu_N$.

Corollary 1 (Vertical integration). *If the supplier fully owns lab 1, then $c = 1 - \mu_N$: the integrated lab pays only the real resource cost. The integrated lab wins with probability $(1 + \mu_N - \alpha) / [2(1 - \alpha)]$, and the independent lab's expected spending falls to $(1 - \alpha - \mu_N)V / [2(1 - \alpha)^2]$.*

This is a form of foreclosure. It does not raise the rival's input price; it makes the integrated lab a stronger competitor, which discourages the rival. Unlike classic foreclosure (Ordover et al., 1990), it reduces total input demand. A supplier that sells to everyone may therefore prefer *not* to integrate, unless $\alpha + \mu_N < 1$.

More generally, a supplier may link to both labs on different terms, so that the labs' normalized costs are $c_1 < c_2 \leq 1$. Lemma 2 then gives total expected compute

$$T(c_1, c_2) = \frac{V}{2(c_2 - \alpha)} + \frac{(c_1 - \alpha)V}{2(c_2 - \alpha)^2} = \frac{(c_1 + c_2 - 2\alpha)V}{2(c_2 - \alpha)^2}. \quad (9)$$

Proposition 4 (Supporting the leader versus the laggard). *For $\alpha < c_1 < c_2$:*

- (a) $\partial T / \partial c_1 = V / [2(c_2 - \alpha)^2] > 0$. *A larger rebate to the stronger lab lowers total compute.*
- (b) $\partial T / \partial c_2 = -(2c_1 + c_2 - 3\alpha)V / [2(c_2 - \alpha)^3] < 0$. *A larger rebate to the weaker lab raises total compute.*
- (c) *Fix the total rebate $r_1 + r_2 = \bar{r}$, with $r_1 \geq r_2$. Total compute is strictly increasing in the smaller rebate r_2 . It is minimized by giving the whole rebate to one lab and maximized by splitting it equally. Hence network value is maximized by the extreme allocation if $\alpha + \mu_N < 1$, and by the equal split if $\alpha + \mu_N > 1$.*

Proof. (a) and (b) follow by differentiating (9). For (c), write $c_1 = 1 - \bar{r} + r_2$ and $c_2 = 1 - r_2$. Then the numerator $c_1 + c_2 - 2\alpha = 2 - \bar{r} - 2\alpha$ does not depend on r_2 , while the denominator $2(1 - r_2 - \alpha)^2$ is decreasing in r_2 . So T is increasing in r_2 on $[0, \bar{r}/2]$. Combine with $\mathbb{E}J = V - (1 - \alpha - \mu_N)T$. $\square$

Proposition 4 explains why the *identity* of the lab a supplier backs matters. Supporting the leading lab strengthens its position and moderates the race; supporting the lagging lab intensifies it. When the race is already too intense from the network's point of view, the supplier should support the leader.




Figure 7: Exclusive versus broad-based investment (Theorem 4), with $\alpha = 0.2$ and $c = 0.8$. The curves cross zero at $\mu_N = 1 - \alpha$ (dotted line).

9 Margins and fragility

Supplier margins in AI compute are high and contested: alternative accelerators, custom silicon and new entrants can compress them. The model shows how circularity interacts with margin risk.

Proposition 5 (Fragility). *Fix holdings (σ, ω, ϕ_L) with $h = 0$ and $\omega > 0$, and let the margin μ vary.*

- (a) *The elasticity of compute spending with respect to the margin is $r/(1 - \alpha - r)$. This is zero without circularity and increasing in r .*
- (b) *The elasticity of the supplier's revenue $2\mu\bar{e}$ with respect to the margin is $(1 - \alpha)/(1 - \alpha - r) > 1$.*
- (c) *The structure's value $\Delta J(\mu) = 2(\alpha + \mu - 1)(\bar{e}(\mu) - \bar{e}^0)$ is positive for $\mu > 1 - \alpha$ and negative for $\mu < 1 - \alpha$. For any $\mu < 1 - \alpha$, $|\Delta J(\mu)|$ is increasing in ω and σ .*

Proof. From (5), $\ln \bar{e} = \text{const} - \ln(1 - \alpha - \omega\mu/\theta)$, so $\partial \ln \bar{e} / \partial \ln \mu = r/(1 - \alpha - r)$, which gives (a). Part (b) adds one. For (c), $\bar{e} - \bar{e}^0 > 0$ is increasing in r , and r is increasing in ω and σ . $\square$

Proposition 5 formalizes a common worry about circular deals. They make compute demand, and with it supplier revenue, more sensitive to supplier margins. A structure designed in a high-margin environment ($\mu > 1 - \alpha$) creates value there, but if margins fall below $1 - \alpha$ the same structure keeps compute high and destroys value. The more circular the structure, the larger the loss. Figure 8 illustrates: with $\sigma = 0.3$ and $\omega = 0.2$, the structure raises compute by 47% and creates value at $\mu = 0.9$; at $\mu = 0.6$ it still raises compute by 27% but destroys value.

Remark 1 (Unwinding). In the model, holdings are fixed when margins change. If unwinding is costly (equity is illiquid, commitments are binding, warrants have vested), the structure persists after margins fall. The result is a period of over-investment that ends only when the structure is unwound. Spending then falls by the factor $(1 - \alpha - r)/(1 - \alpha)$, and the supplier's revenue falls by more than the margin alone would imply.




Figure 8: Fragility (Proposition 5). Left: a fixed circular structure creates value only when $\mu > 1 - \alpha = 0.8$ (dotted line). Right: circularity raises the sensitivity of compute and of supplier revenue to the margin.

10 Designing circular deals

So far, holdings were taken as given. Suppose instead that the insiders choose them, together with cash transfers. In practice the transfers are the cash a supplier pays for its stake, or the price at which warrants are granted. Insiders' control imposes limits: lab insiders must retain at least $\underline{\theta}$ of their lab, and the supplier insider must retain at least $\underline{s}$ of the supplier, so $\omega \leq \bar{\omega} \equiv (1 - \phi_S - \underline{s})/2$. To keep compute finite, assume $\bar{\omega}\mu < \underline{\theta}(1 - \alpha)$.

Proposition 6 (Optimal circularity is all or nothing). *Consider symmetric (inclusive) structures, chosen to maximize the insiders' combined payoff with transfers. Let*

$$\Lambda \equiv \mu(1 - \phi_S) - (1 - \alpha)(1 - \phi_L).$$

- (a) *The insiders' combined payoff is affine in spending per lab, with slope 2Λ , for every structure (σ, ω, h) .*
- (b) *If $\Lambda > 0$, the optimal structure maximizes the rebate: $h = 0$, $\omega = \bar{\omega}$ and $\theta = \underline{\theta}$ (so $\sigma = 1 - \underline{\theta} - \phi_L$). Compute is $\bar{e}^* = V/[2(1 - \alpha - \bar{\omega}\mu/\underline{\theta})]$.*
- (c) *If $\Lambda < 0$, the optimal structure has no reverse stake ($\omega = 0$) and maximizes horizontal holdings h : the insiders want less compute.*
- (d) *With Nash bargaining between the supplier insider and the lab insiders, each side's gain equals half of $2\Lambda(\bar{e}^* - \bar{e}^0)$ (case b), net of the transfer that implements the split.*

Proof. (a) Firm profits depend on the structure only through $\bar{e}$: $\mathbb{E}\Pi_i = V/2 - (1 - \alpha)\bar{e}$ and $\mathbb{E}\Pi_S = 2\mu\bar{e}$. Hence $J_{\text{ins}} = \text{const} + 2\bar{e}\Lambda$. (b), (c) By (5), $\bar{e}$ is increasing in $r = \omega\mu/\theta$, decreasing in h , and has its supremum at the corner stated. (d) Transfers are lump sum, so Nash bargaining splits the surplus equally. $\square$

The proposition makes a clear prediction: circular structures should either be absent or extend to the limits that control and governance allow. Since Λ falls with margin compression and rises with outside financing of labs, the model predicts waves: circular deals appear when supplier margins are high and lab funding is abundant, and disappear when margins compress.

The supplier's cash. In the model, the stake the supplier acquires is worth the labs' expected profits, which become negative once circularity raises spending (Proposition 2). A supplier that pays cash for such a stake is not buying a profitable claim. It is paying for the increase in its own margin revenue that the structure induces. Part of the supplier's revenue is thus funded by its own equity investment, and the rest by the labs' outside investors. This is the precise sense in which the arrangement is "circular".

11 Many labs

The two-lab race keeps the algebra simple. The main formulas extend to $n \geq 2$ labs. Let the winner earn $V + \alpha \sum_{k=1}^n e_k$, let each lab insider hold θ of its lab and ω of the supplier, and let the supplier insider hold σ of each lab ($h = 0$).

Proposition 7 (Many labs). *In any symmetric equilibrium in which spending is atomless and zero is in the support, expected spending per lab and total spending are*

$$\bar{e}_n = \frac{V}{n(1 - \alpha - r)}, \quad T_n = n\bar{e}_n = \frac{V}{1 - \alpha - r}, \quad r = \frac{\omega\mu}{\theta}.$$

The network-value identity $\mathbb{E}J = V - (1 - \alpha - \mu_N)T$ and the threshold $\alpha + \mu_N \geq 1$ continue to hold.

Proof. Each lab wins with probability $1/n$, so $\mathbb{E}\Pi_i = (V + \alpha n\bar{e}_n)/n - \bar{e}_n$ and $\mathbb{E}\Pi_S = \mu n\bar{e}_n$. A lab insider's equilibrium payoff equals the payoff from spending zero, which loses with probability one: $\omega\mu(n - 1)\bar{e}_n$. Equating with $\theta\mathbb{E}\Pi_i + \omega\mathbb{E}\Pi_S$ gives $\theta V/n = \bar{e}_n[\theta(1 - \alpha) - \omega\mu]$. The value identity follows as in Theorem 2. $\square$

Total compute does not depend on the number of labs. Entry into the frontier race reallocates spending but does not raise it, because rents are fully dissipated whatever the number of contestants. The circularity multiplier $(1 - \alpha)/(1 - \alpha - r)$ is also independent of n . It is circularity, not entry, that determines the scale of compute spending.

12 Financing constraints: when one-way stakes matter

Proposition 1(a) relies on labs being able to finance any level of spending. Suppose instead that each lab can spend at most its cash B . Take $\alpha = 0$ to keep the algebra simple, and let insiders' effective cost be $c = 1 - r$ per dollar.

Proposition 8 (Budget caps). *With budget B per lab:*

- (a) *If $B \geq V/c$, the cap never binds. If $V/(2c) \leq B < V/c$, there is an equilibrium in which each lab mixes uniformly with density c/V on $[0, 2B - V/c]$ and places an atom of mass $2(1 - cB/V)$ at B . Expected spending is $V/(2c)$, as without the cap, and rents are zero.*
- (b) *If $B < V/(2c)$, both labs spend B with probability one, and each lab insider earns $V/2 - cB > 0$ from the race.*

Hence spending per lab is $\min\{B, V/(2(1 - r))\}$.

Proof. (a) A lab spending $e < 2B - V/c$ wins with probability ce/V , so its payoff is $V \cdot ce/V - ce = 0$. Spending B wins with probability $1 - q + q/2$ with $q = 2(1 - cB/V)$, which gives $V(1 - q/2) - cB = 0$. Spending between $2B - V/c$ and B wins with probability $1 - q$ at strictly higher cost, so it gives a negative payoff. Expected spending equals $V/(2c)$ because rents are zero and $\mathbb{E}[\text{prize}] = V/2$ per lab. (b) At the profile (B, B) , a lab's payoff is $V/2 - cB > 0$; spending less loses for sure, and spending more is infeasible. $\square$

With financing constraints, the two sides of a circular deal play different roles. The supplier's cash (the stake σ) relaxes the cap B and raises spending one-for-one, but only up to the unconstrained level $V/(2(1 - r))$. The reverse stake raises that level itself. Cash helps only when financing binds, whereas rebates raise the level of spending that labs would choose without constraints. Once labs are well funded, as frontier labs largely are, only the rebate matters, and Proposition 1(a) applies. The value test of Theorem 2 is unchanged, because it is an accounting identity.

13 Data centers and power: a two-layer stack

So far, the supplier simply collects a margin. But data centers and the power to run them must be built, and that buildout is itself competitive and subject to spillovers. This section adds a second racing layer, bringing the model closer to the two-industry structure of Xiong (2026).

13.1 Setup

There are two *infrastructure* firms, P_1 and P_2 (data-center or power developers), which race to build capacity $k_g \geq 0$ for an anchor contract. The two layers are linked as follows:

- The lab race has prize $R_L = V_L + \alpha(e_1 + e_2) + \beta_L(k_1 + k_2)$: capacity raises the value of frontier models, for instance by lowering inference costs.
- The infrastructure race has prize $R_P = V_P + \alpha_P(k_1 + k_2) + \beta_P(e_1 + e_2)$: lab compute raises the value of capacity contracts, and α_P captures learning and shared grid upgrades.

- A share μ of lab spending is profit for the infrastructure layer, split equally between P_1 and P_2 (labs multi-home).

Ownership: each lab insider holds θ_L of its lab, h of the rival lab and ω of each infrastructure firm. Each infrastructure insider holds θ_P of its firm, h_P of the rival infrastructure firm and σ of each lab. Closedness requires $\theta_L = 1 - h - 2\sigma$ and $\theta_P = 1 - h_P - 2\omega$. Firm profits are $\Pi_i = R_L \mathbf{1}\{i \text{ wins}\} - e_i$ for labs and $\Pi_g = R_P \mathbf{1}\{g \text{ wins}\} - k_g + \frac{\mu}{2}(e_1 + e_2)$ for infrastructure firms.

13.2 Equilibrium and value

Theorem 5 (Two-layer stack). *Let $m_L = \mathbb{E}e_i$ and $m_P = \mathbb{E}k_g$. Suppose $\theta_L > h$, $\theta_P > h_P$, $\alpha, \alpha_P > 0$,*

$$\kappa_L - d_L \equiv \theta_L(1 - \alpha) - \omega(\beta_P + \mu) > 0, \quad \kappa_P - d_P \equiv \theta_P(1 - \alpha_P) - \sigma\beta_L > 0,$$

and let

$$a_L = \frac{\theta_L - h}{2[\theta_L(1 - \alpha) - \omega(\beta_P + \mu)]}, \quad a_P = \frac{\theta_P - h_P}{2[\theta_P(1 - \alpha_P) - \sigma\beta_L]},$$

with $4a_L a_P \beta_L \beta_P < 1$ and positive intercepts. Then both races have equilibria of the form (3), and

$$m_L = \frac{a_L(V_L + 2\beta_L a_P V_P)}{1 - 4a_L a_P \beta_L \beta_P}, \quad m_P = \frac{a_P(V_P + 2\beta_P a_L V_L)}{1 - 4a_L a_P \beta_L \beta_P}. \quad (10)$$

Network value is

$$\mathbb{E}J = V_L + V_P - 2(1 - \alpha - \beta_P - \mu)m_L - 2(1 - \alpha_P - \beta_L)m_P. \quad (11)$$

Proof. For a lab insider, the win-minus-lose payoff is $(\theta_L - h)R_L$. Expected infrastructure capacity enters only through $\mathbb{E}(k_1 + k_2) = 2m_P$, so $\Delta = (\theta_L - h)(V_L + 2\beta_L m_P)$ and $d = (\theta_L - h)\alpha$. The insider's stake ω in each infrastructure firm adds $\omega(\beta_P + \mu)$ per dollar of own spending, because $\Pi_{P_1} + \Pi_{P_2} = R_P - k_1 - k_2 + \mu(e_1 + e_2)$. Hence $\kappa = \theta_L - h\alpha - \omega(\beta_P + \mu)$, and Lemma 1 gives $m_L = a_L(V_L + 2\beta_L m_P)$. Symmetrically, $m_P = a_P(V_P + 2\beta_P m_L)$, since an infrastructure insider's stake σ in each lab adds $\sigma\beta_L$ per unit of own capacity. Solving the linear system gives (10). Summing firm profits, and using that exactly one firm wins each race, gives (11). $\square$

Corollary 2 (Value thresholds with cross-layer feedback). *Start from the benchmark ($h = h_P = \sigma = \omega = 0$).*

- (a) Labs invest in infrastructure ($\omega > 0, \sigma = 0$). If $\beta_P + \mu > 0$, lab compute rises, and network value rises if and only if

$$\alpha + \mu + \frac{\beta_L \beta_P}{1 - \alpha_P} > 1.$$

- (b) Infrastructure invests in labs ($\sigma > 0, \omega = 0$). If $\beta_L > 0$, capacity rises, and network value rises if




Figure 9: Corollary 2(a) with $\alpha = \alpha_P = 0.2$ and $\beta_L = 0.3$. Cross-layer feedback lowers the margin needed for labs' stakes in infrastructure to create value. With these parameters, infrastructure's stakes in labs create value only if $\mu > (1 - \alpha_P)(1 - \alpha) / \beta_L - \beta_P \geq 1.53$, which is infeasible.

and only if

$$\alpha_P + \frac{\beta_L(\mu + \beta_P)}{1 - \alpha} > 1.$$

Proof. For (a): with $\sigma = 0$, $a_P = 1/[2(1 - \alpha_P)]$ does not depend on ω , so $\Delta m_P = 2a_P\beta_P\Delta m_L$ exactly. Substituting into (11),

$$\Delta J = -2\Delta m_L \left[(1 - \alpha - \beta_P - \mu) + \frac{\beta_P(1 - \alpha_P - \beta_L)}{1 - \alpha_P} \right] = -2\Delta m_L \left[1 - \alpha - \mu - \frac{\beta_L\beta_P}{1 - \alpha_P} \right].$$

Part (b) is symmetric, with $\Delta m_L = 2a_L\beta_L\Delta m_P$ and $a_L = 1/[2(1 - \alpha)]$. $\square$

Corollary 2 has three implications. First, the threshold for labs' stakes in infrastructure (think of stakes in data-center developers) is the single-stack threshold $\alpha + \mu$ plus a *feedback term* $\beta_L\beta_P/(1 - \alpha_P)$: more lab compute raises the value of capacity, more capacity raises the value of the frontier, and the loop is amplified by the infrastructure layer's own spillover. Second, infrastructure firms' stakes in labs pay off only if capacity spills over strongly onto the frontier (β_L large), and then they are valued *through* the margin and demand linkages $\mu + \beta_P$. Third, without margins ($\mu = 0$) and with all spillovers below one half, both thresholds fail, and the impossibility theorem of Xiong (2026) reappears as a special case.

13.3 Power and leakage

Power is often supplied by independent or regulated producers outside the ownership network. Their margin μ_O is a transfer from the network's perspective but a *leakage*: it is paid by labs and not recovered by insiders.

Proposition 9 (Bringing power into the network). *Suppose lab insiders acquire stakes ω_O in the outside layer (for instance through equity in power projects) in addition to ω in the supplier. Then the effective rebate*

becomes $r = (\omega\mu + \omega_O\mu_O)/\theta$; compute is given by (5) with this r ; and the value threshold of Theorem 2 becomes $\alpha + \mu + \mu_O > 1$. Integrating power into the network can therefore turn compute-raising circular structures from value-destroying into value-creating.

Proof. Lab insider i 's payoff gains $\omega_O\mu_O(e_i + e_j)$, so κ falls by $\omega_O\mu_O$. Network profits now include Π_O , so the in-network margin is $\mu + \mu_O$ in (7). $\square$

The in-network margin μ_N is thus the key state variable. It is not a property of any single firm; it is the share of a compute dollar that is captured *somewhere* in the ownership network.

14 Welfare and policy

14.1 Private versus social thresholds

Let η be the external value of a dollar of compute to parties outside the industry. It is positive if frontier progress diffuses to users and other sectors; negative if compute imposes grid congestion, emissions or safety risk. A planner counts all margins as transfers. With total margin $\mu_{\text{tot}} = \mu_N + \mu_{\text{out}}$ (in-network plus outside layers), social welfare changes by

$$\Delta SW = (\alpha + \eta + \mu_{\text{tot}} - 1) \Delta T. \quad (12)$$

Proposition 10 (Private versus social value of circularity). *Assume a closed network.*

- (a) *If $\eta \geq 0$, every inclusive circular structure that is privately profitable ($\alpha + \mu_N > 1$) is socially beneficial. There are also socially beneficial structures that are not privately profitable, whenever $\alpha + \mu_N < 1 < \alpha + \eta + \mu_{\text{tot}}$.*
- (b) *If $\eta \geq 0$, exclusive deals and vertical integration that are privately profitable ($\alpha + \mu_N < 1$) are socially harmful whenever $\alpha + \eta + \mu_{\text{tot}} > 1$.*
- (c) *If $\eta < 0$, privately profitable inclusive structures are socially harmful whenever $\alpha + \eta + \mu_{\text{tot}} < 1 < \alpha + \mu_N$.*

Proof. Inclusive structures raise T and exclusive structures lower it (Theorems 1 and 4). Compare the signs of $\alpha + \mu_N - 1$ and $\alpha + \eta + \mu_{\text{tot}} - 1$, using $\mu_{\text{tot}} \geq \mu_N$. $\square$

14.2 Policy implications

Antitrust. Exclusive circular deals and vertical integration lower total compute and raise the favored lab's probability of reaching the frontier first (Theorem 4, Corollary 1). They are privately profitable exactly when inclusive deals are not, and they can harm welfare when compute has positive external value. This is a novel foreclosure channel: discouragement in an innovation race

rather than raising rivals' costs. By contrast, inclusive circular structures raise compute and are more likely to be benign.

Disclosure and investor protection. Theorem 3 identifies the labs' outside investors as the party that bears the losses when circular structures raise compute beyond its value to the network. Disclosure of reverse stakes, warrants that vest with purchases, and commitments linked to investments would let outside investors price this transfer.

Energy. Bringing power producers into the ownership network raises compute (Proposition 9). Whether that is desirable depends on η , which for energy-intensive compute includes grid and emissions externalities that the network does not internalize.

Product-market conduct. I have assumed that ownership affects only investment. If cross-holdings also soften competition between the labs once a frontier is reached, there is an additional, conventional anticompetitive effect, as in Reynolds and Snapp (1986). The horizontal parameter h captures the investment side of that effect: it lowers compute (Proposition 1).

15 An illustrative calibration

The model's main thresholds involve two numbers: the knowledge spillover α and the in-network margin μ_N , the share of a lab's compute dollar that becomes profit inside its ownership network. Table 1 computes μ_N for hypothetical cost shares and margins. The numbers are *illustrative assumptions*, not estimates; they show how the thresholds would be applied.

Layer	Share of compute dollar	Margin on that share	Contribution to μ
Accelerators	0.45	0.65	0.29
Data center and cloud	0.40	0.30	0.12
Power	0.15	0.20	0.03

Network contains	μ_N	$\alpha + \mu_N$ ($\alpha = 0.1$)	$\alpha + \mu_N$ ($\alpha = 0.3$)
Accelerator supplier only	0.29	0.39	0.59
Accelerators and cloud	0.41	0.51	0.71
Accelerators, cloud and power	0.44	0.54	0.74

Table 1: Illustrative in-network margins. All entries are hypothetical.

Under these assumptions, $\alpha + \mu_N < 1$ in every case. The race over-invests from the network's point of view, and Theorem 2(a) applies: closed-network circular structures that raise compute destroy value, while exclusive deals create it. Two qualifications matter. First, by Theorem 3,

compute-raising structures are nevertheless profitable for insiders when the labs' outside share is large enough. With $\phi_S = 0$, $\alpha = 0.3$ and $\mu_N = 0.41$, the condition is $\phi_L > 1 - 0.41/0.7 \approx 0.41$, a plausible level for labs funded by repeated outside rounds. Second, α is hard to measure: diffusion through distillation, open weights and staff mobility may be large, and Corollary 2 shows that cross-layer feedback adds to it. The regime boundary $\alpha + \mu_N = 1$ is the object that empirical work should target.

16 Discussion of modeling choices

Why a contest? Frontier AI has strong winner-take-most features: being first to a capability level brings users, talent and pricing power, while followers face commoditized markets. An all-pay contest captures this and the fact that compute is sunk whether or not the lab wins. The contest also delivers full rent dissipation in the benchmark, which makes every effect of ownership visible in a single number, $\bar{e}$. In Xiong (2026), Cournot-type models of R&D give different thresholds, because R&D there creates joint surplus. In the present setting, the analogue of that surplus is the margin μ and the cross-layer spillovers, which is why the thresholds here can go either way.

Exogenous margins. Margins are taken as given, and Section 9 studies changes in them. If the supplier set its price, and lab spending were measured in dollars, a monopolist would raise its price without bound in this model: with full dissipation, labs spend the same dollars whatever the price. Margins should therefore be read as determined by upstream competition and capacity, which lie outside the model.

Passive outsiders and insider control. Insiders control decisions while holding minority economic stakes, as is common for founder-led labs and for firms with dual-class shares. This separation is what allows outside finance to amplify rebates (Proposition 1(c)) and to absorb losses (Theorem 3).

Linearity. Prizes and margins are linear in spending. This yields closed forms, and it makes the value identity (7) exact. With decreasing returns, the thresholds become local conditions at the equilibrium, as in Theorem 1 of Xiong (2026).

17 Testable implications

The model yields predictions that distinguish its mechanisms from alternative explanations of circular deals, such as financing constraints or risk sharing.

1. **Reverse stakes, not supplier stakes, move compute** (Proposition 1). A supplier's equity in a lab should not, by itself, raise the lab's compute purchases once financing is controlled for.

Instruments that pay the lab per unit purchased (warrants vesting with deployment, equity in the supplier, volume rebates) should. A financing-constraint explanation predicts the opposite ranking.

2. **Amplification by outside finance** (Propositions 1(c) and 6). Holding the reverse stake fixed, labs whose insiders own a smaller share should respond more strongly to it.
3. **Margin sensitivity** (Proposition 5). The elasticity of a lab’s compute spending with respect to its supplier’s margin, for example around price cuts or the entry of competing accelerators, should increase with the lab’s effective rebate. Suppliers’ revenue should react more than one-for-one to margin shocks when their customers are circularly linked.
4. **Accounting losses and outside returns** (Proposition 2, Theorem 3). Labs in more circular structures should report larger operating losses, and their outside investors should earn lower returns, while insiders’ total compensation, including claims on suppliers, need not fall.
5. **Exclusive ties discourage rivals** (Theorem 4). After an exclusive deal, the favored lab’s spending should not rise, while rival labs’ spending should fall and the favored lab should lead more often.
6. **Breadth of investment** (Theorem 4). Suppliers with higher margins on lab spending should be more likely to invest across many labs; lower-margin or vertically integrated investors should concentrate on one.
7. **Power in the network** (Proposition 9). Equity links between AI firms and power producers should be followed by higher compute spending than comparable contract-only arrangements.

18 Conclusion

Circular investment links AI labs to the firms that sell them compute. This paper models the AI stack as a rent-dissipating frontier race layered on top of a margin-earning compute and infrastructure sector, with cross-holdings among them. One-way stakes are neutral; reverse stakes act as volume rebates, and the supplier’s stakes amplify them. Whether the resulting increase in compute creates value depends on a single threshold, $\alpha + \mu_N \gtrless 1$, comparing the knowledge spillover plus the in-network margin with the dollar paid. Below the threshold, a circular-investment impossibility theorem holds; above it, circularity corrects a double-marginalization problem in compute. The labs’ outside investors bear the losses when the increase destroys value. Exclusive deals and inclusive deals are mirror images. Circularity amplifies margin risk. With data centers and power, cross-layer spillovers add feedback terms to the thresholds, and the boundary of the ownership network determines which margins count.

Several extensions are natural: endogenous margins set by competing suppliers, financing constraints (under which the supplier’s cash can matter even without reverse stakes), more than

two labs, dynamic races, and estimation of α and μ_N .

A Proofs

Proof of Lemma 2. Normalize payoffs so that player p earns $V + \alpha(e_1 + e_2) - c_p e_p$ if it wins and $-c_p e_p$ if it loses. The weak player's indifference against the strong player's c.d.f. G_1 requires, for e in the support,

$$U_2(e) = \int_{[0,e)} (V + \alpha e + \alpha y) dG_1(y) - c_2 e = \text{const.}$$

Differentiating gives $g_1(e)(V + 2\alpha e) + \alpha G_1(e) = c_2$. Solving as in Lemma 1 gives $G_1\sqrt{u} = (c_2/\alpha)\sqrt{u} + C$, and $G_1(0) = 0$ yields $G_1 = (c_2/\alpha)[1 - \sqrt{V/u}]$. The strong player's indifference against G_2 , which may have an atom at 0, gives the same equation with c_1 (the atom contributes $\alpha G_2(0)$ to the derivative, which is included in $\alpha G_2(e)$). So $G_2 = c_1/\alpha + C'/\sqrt{u}$. Both distributions reach 1 at the common upper bound b : $G_1(b) = 1$ gives $\sqrt{V/u_b} = 1 - \alpha/c_2$, and $G_2(b) = 1$ then gives $C' = -(c_1 - \alpha)c_2\sqrt{V}/[\alpha(c_2 - \alpha)]$. G_2 is increasing because $C' < 0$, and $G_2(0) = (c_2 - c_1)/(c_2 - \alpha) \in (0, 1)$.

The weak player earns $U_2(0) = 0$, and by construction U_2 is constant on $[0, b]$. The strong player never bids exactly 0: bidding 0^+ beats the atom and earns $V G_2(0)$, more than the tie payoff at 0. Hence $U_1 \equiv V G_2(0)$ on $(0, b]$. Above b , both payoffs are strictly decreasing, since $c_p > \alpha$.

For the means, write $\mathbb{E}e = \int_0^b (1 - G) de$, with $de = du/(2\alpha)$ and $\int u^{-1/2} du = 2\sqrt{u}$. For G_1 , the computation is that of Lemma 1 with $\kappa/d = c_2/\alpha$, giving $V/[2(c_2 - \alpha)]$. For G_2 ,

$$\mathbb{E}e_2 = \left(1 - \frac{c_1}{\alpha}\right)b + \frac{(c_1 - \alpha)c_2\sqrt{V}}{\alpha^2(c_2 - \alpha)}(\sqrt{u_b} - \sqrt{V}).$$

Here $b = V(2c_2 - \alpha)/[2(c_2 - \alpha)^2]$ and $\sqrt{u_b} - \sqrt{V} = \alpha\sqrt{V}/(c_2 - \alpha)$, so $\mathbb{E}e_2 = (c_1 - \alpha)V/[2(c_2 - \alpha)^2]$. Finally,

$$\Pr(e_1 > e_2) = \int G_2 dG_1 = \frac{c_2\sqrt{V}}{2\alpha} \int_V^{u_b} \left(\frac{c_1}{\alpha} - \frac{(c_1 - \alpha)c_2\sqrt{V}}{\alpha(c_2 - \alpha)}u^{-1/2}\right)u^{-3/2}du,$$

which simplifies to $(2c_2 - c_1 - \alpha)/[2(c_2 - \alpha)]$. □

Proof of Corollary 1. With full ownership, the integrated firm's payoff from the lab's spending is $R1\{\text{win}\} - e_1 + \mu_N(e_1 + e_2)$, so its normalized marginal cost is $1 - \mu_N$. Apply Theorem 4(a) with $c = 1 - \mu_N$. □

B Numerical verification

All formulas were verified numerically.

- *Lemma 1.* For 30 random coefficient vectors, the payoff against F is constant on the support (maximum deviation 3×10^{-4} , attributable to the integration grid) and strictly lower above it; the numerical mean matches (4).
- *Theorem 1 and Theorem 2.* Over 20,000 random parameter draws, (6) matches the portfolio payoff and (7) holds, both to machine precision (2×10^{-15}).
- *Lemma 2.* For $\alpha = 0.25$ and $c_1 = 0.7$: the atom is 0.400 (formula 0.400); the strong player's payoff varies by less than 10^{-6} on its support; the means are 0.6667 and 0.4000 (formulas 0.6667 and 0.4000); the win probability is 0.7000 (formula 0.7000); and both payoffs fall above the support.
- *Corollary 2.* In 19,770 random parameter draws, the sign of ΔJ for one-sided deals agrees with the stated thresholds in every case.

References

- AMD (2025). AMD and OpenAI announce strategic partnership to deploy 6 gigawatts of AMD GPUs. Press release, October 6, 2025. <https://newsroom.amd.com/news/amd-and-openai-announce-strategic-partnership-to-d/>
- Antón, M., Ederer, F., Giné, M., & Schmalz, M. (2021). Innovation: The bright side of common ownership? Working paper.
- Baye, M. R., Kovenock, D., & de Vries, C. G. (1996). The all-pay auction with complete information. *Economic Theory*, 8(2), 291–305.
- Baye, M. R., Kovenock, D., & de Vries, C. G. (2012). Contests with rank-order spillovers. *Economic Theory*, 51(2), 315–350.
- Bloomberg (2026). Nvidia invests \$2 billion more in CoreWeave, offers new chip. January 26, 2026. <https://www.bloomberg.com/news/articles/2026-01-26/nvidia-invests-another-2-billion-in-coreweave-offers-new-chip>
- Brennan, M. J., Maksimovic, V., & Zechner, J. (1988). Vendor financing. *Journal of Finance*, 43(5), 1127–1141.
- Business Standard (2025). Microsoft retains 27% stake in OpenAI valued at \$135 bn in new deal. October 28, 2025. https://www.business-standard.com/world-news/microsoft-retains-27pc-stake-in-openai-valued-at-135-bn-new-deal-125102801567_1.html
- Carnegie Endowment for International Peace (2026). Beyond the hype: Assessing hyperscaler nuclear commitments against U.S. energy re-

- alities. June 2026. <https://carnegieendowment.org/research/2026/06/beyond-the-hype-assessing-hyperscaler-nuclear-commitments-against-us-energy-realities>
- d'Aspremont, C., & Jacquemin, A. (1988). Cooperative and noncooperative R&D in duopoly with spillovers. *American Economic Review*, 78(5), 1133–1137.
- Flath, D. (1989). Vertical integration by means of shareholding interlocks. *International Journal of Industrial Organization*, 7(3), 369–380.
- Gilo, D., Moshe, Y., & Spiegel, Y. (2006). Partial cross ownership and tacit collusion. *RAND Journal of Economics*, 37(1), 81–99.
- Greenlee, P., & Raskovich, A. (2006). Partial vertical ownership. *European Economic Review*, 50(4), 1017–1041.
- Hillman, A. L., & Riley, J. G. (1989). Politically contestable rents and transfers. *Economics & Politics*, 1(1), 17–39.
- Kamien, M. I., Muller, E., & Zang, I. (1992). Research joint ventures and R&D cartels. *American Economic Review*, 82(5), 1293–1306.
- Levy, N., Spiegel, Y., & Gilo, D. (2018). Partial vertical integration, ownership structure, and foreclosure. *American Economic Journal: Microeconomics*, 10(1), 132–180.
- López, Á. L., & Vives, X. (2019). Overlapping ownership, R&D spillovers, and antitrust policy. *Journal of Political Economy*, 127(5), 2394–2437.
- Morningstar (2025). What are circular AI chip deals, and should investors be worried? <https://global.morningstar.com/en-gb/stocks/what-are-circular-ai-chip-deals-should-investors-be-worried>
- Ordover, J. A., Saloner, G., & Salop, S. C. (1990). Equilibrium vertical foreclosure. *American Economic Review*, 80(1), 127–142.
- Pure AI (2025). Microsoft, Nvidia deepen ties with Anthropic as startup commits \$30 billion to Azure. November 18, 2025. <https://pureai.com/Articles/2025/11/18/Microsoft-and-Nvidia-deepen-ties-with-Anthropic.aspx>
- Reynolds, R. J., & Snapp, B. R. (1986). The competitive effects of partial equity interests and joint ventures. *International Journal of Industrial Organization*, 4(2), 141–153.
- Siegel, R. (2009). All-pay contests. *Econometrica*, 77(1), 71–92.
- The Next Web (2026). NVIDIA tops \$40bn in AI equity bets in 2026, anchored by \$30bn OpenAI investment. <https://thenextweb.com/news/nvidia-40bn-ai-equity-investments-2026>

The Register (2025). AI's trillion dollar deal wheel bubbling around Nvidia, OpenAI. November 4, 2025. https://www.theregister.com/2025/11/04/the_circular_economy_of_ai/

Xiong, N. (2026). R&D contests with spillovers and cross shareholdings across industries. Working paper.