

Additional Materials to “The Macroeconomic Effect of AI: Sizing the Software Engineering Channel”¹

Alex Blumenfeld Jonathon Hazell Chen Lian Andreas Schaab

This document contains additional material for the article “[The Macroeconomic Effect of AI: Sizing the Software Engineering Channel](#).” References to equations or sections without a capital-letter prefix refer to the main article; those with prefixes A–D refer to the [supplementary appendix](#).

¹Blumenfeld: UC Berkeley, alex.blumenfeld@berkeley.edu; Hazell: London School of Economics, j.hazell@lse.ac.uk; Lian: UC Berkeley and NBER, chen_lian@berkeley.edu; Schaab: UC Berkeley and NBER, schaab@berkeley.edu.

E Additional Tables

Table E.1: Twenty largest holdings in the THNQ ETF

Rank	Company	Portfolio Weight (%)
1	Lumentum Holdings (LITE)	5.05
2	Palo Alto Networks (PANW)	3.95
3	Lam Research (LRCX)	3.12
4	Nebius Group (NBIS)	2.97
5	TSMC (2330.TW)	2.97
6	Teradyne (TER)	2.84
7	ASML Holding (ASML)	2.79
8	Analog Devices (ADI)	2.65
9	Cloudflare (NET)	2.64
10	Cognex (CGNX)	2.57
11	MediaTek (2454.TW)	2.50
12	Alphabet (GOOGL)	2.44
13	Global Unichip (3443.TW)	2.38
14	Arista Networks (ANET)	2.27
15	NVIDIA (NVDA)	2.25
16	Infineon Technologies (IFX)	2.25
17	Advanced Micro Devices (AMD)	2.24
18	Meta Platforms (META)	2.21
19	Amazon.com (AMZN)	2.21
20	Pure Storage (PSTG)	2.16
Top 20 Total		54.46

Notes: This table reports the twenty largest holdings in the THNQ ETF by portfolio weight as of March 21, 2026. The fund held 51 securities in total on this date. Source: ROBO Global ETFs (<https://www.roboglobaletf.com/thnq>).

Table E.2: Twenty largest holdings in the AIQ ETF

Rank	Company	Portfolio Weight (%)
1	SK hynix (000660.KRX)	4.34
2	Samsung Electronics (005930.KRX)	4.26
3	Netflix (NFLX)	3.65
4	Micron Technology (MU)	3.54
5	Cisco Systems (CSCO)	3.47
6	TSMC (TSM)	3.34
7	Apple (AAPL)	3.32
8	Broadcom (AVGO)	3.16
9	NVIDIA (NVDA)	3.11
10	Meta Platforms (META)	3.05
11	Alphabet (GOOGL)	3.03
12	Palantir Technologies (PLTR)	3.02
13	Amazon.com (AMZN)	2.91
14	Microsoft (MSFT)	2.81
15	Tencent Holdings (0700.HK)	2.81
16	Oracle (ORCL)	2.79
17	Tesla (TSLA)	2.72
18	AMD (AMD)	2.64
19	IBM (IBM)	2.33
20	Alibaba Group (BABA)	2.30
Top 20 Total		62.60

Notes: This table reports the twenty largest holdings in the AIQ ETF by portfolio weight as of March 21, 2026. The fund held 84 securities in total on this date. Source: Global X ETFs (<https://www.globalxetfs.com/funds/aiq>).

Table E.3: GPT-5.6-classified semiconductor-related firms

<i>Panel A: NAICS6 composition of GPT-5.6-classified semiconductor-related firms</i>		
NAICS		Count of firms
334413: Semiconductor and Related Device Manufacturing		39
333242: Semiconductor Machinery Manufacturing		9
334513: Instruments and Related Products Manufacturing for Measuring, Displaying, and Controlling Industrial Process Variables		3
334515: Instrument Manufacturing for Measuring and Testing Electricity and Electrical Signals		3
513210: Software Publishers		3
325120: Industrial Gas Manufacturing		1
331410: Nonferrous Metal (except Aluminum) Smelting and Refining		1
333310: Commercial and Service Industry Machinery Manufacturing		1
333314: Missing / Unknown		1
334511: Search, Detection, Navigation, Guidance, Aeronautical, and Nautical System and Instrument Manufacturing		1
336110: Automobile and Light Duty Motor Vehicle Manufacturing		1
Missing: Missing / Unknown		1
<i>Panel B: 25 largest GPT-5.6-classified semiconductor-related firms</i>		
Rank	Firm	Market Cap (2021)
1	LINDE PLC	\$177.6 B
2	TEXAS INSTRUMENTS INC	\$174.1 B
3	APPLIED MATERIALS INC	\$139.5 B
4	KLA CORP	\$64.8 B
5	MICROCHIP TECHNOLOGY INC	\$48.4 B
6	STMICROELECTRONICS NV	\$44.5 B
7	GLOBALFOUNDRIES INC	\$34.6 B
8	ON SEMICONDUCTOR CORP	\$29.4 B
9	SKYWORKS SOLUTIONS INC	\$25.5 B
10	MONOLITHIC POWER SYSTEMS INC	\$22.8 B

Continued on next page

Table E.3: GPT-5.6-classified semiconductor-related firms (continued)

Panel B continued

Rank	Firm	Market Cap (2021)
11	ENTEGRIS INC	\$18.8 B
12	QORVO INC	\$17.0 B
13	WOLFSPEED INC	\$13.0 B
14	SYNAPTICS INC	\$11.4 B
15	LATTICE SEMICONDUCTOR CORP	\$10.6 B
16	MKS INC	\$9.7 B
17	IPG PHOTONICS CORP	\$9.1 B
18	SILICON LABORATORIES INC	\$7.9 B
19	SITIME CORP	\$6.1 B
20	AMKOR TECHNOLOGY INC	\$6.1 B
21	MAXLINEAR INC	\$5.8 B
22	SEMTECH CORP	\$5.7 B
23	POWER INTEGRATIONS INC	\$5.6 B
24	MACOM TECHNOLGY SOL HLDGS INC	\$5.5 B
25	CIRRUS LOGIC INC	\$5.3 B

Notes: Both panels restrict attention to firms that remain after excluding THNQ/AIQ constituents and before dropping semiconductor-related firms; this yields 64 firms. Panel A reports the NAICS6 composition of firms that GPT 5.6 classified as semiconductor-related. Panel B reports the 25 largest GPT-5.6-classified semiconductor-related firms ranked by end-of-2021 market capitalization. Panel A reports NAICS entries as code followed by description. Market capitalization is reported in billions of U.S. dollars.

Table E.4: Firms with the highest and lowest software engineer payroll shares

Rank	Firm	NAICS2	Market Cap (2021)
<i>Panel A: 20 firms with the highest software engineer payroll shares</i>			
1	CIMPRESS PLC	31-33	\$1.9 B
2	INNOVIZ TECHNOLOGIES LTD	31-33	\$0.9 B
3	UBIQUITI INC	31-33	\$19.2 B
4	MASTECH DIGITAL INC	56	\$0.2 B
5	CIENA CORP	31-33	\$11.9 B
6	ENERGOUS CORP	31-33	\$0.1 B
7	ARLO TECHNOLOGIES INC	31-33	\$0.9 B
8	PAYPAL HOLDINGS INC	52	\$220.3 B
9	MARQETA INC	52	\$7.2 B
10	GARMIN LTD	31-33	\$26.2 B
11	SHIFT4 PAYMENTS INC	52	\$3.0 B
12	SONOS INC	31-33	\$3.8 B
13	AEVA TECHNOLOGIES INC	31-33	\$1.6 B
14	CAMBIUM NETWORKS CORP	31-33	\$0.7 B
15	EXPEDIA GROUP INC	56	\$27.1 B
16	BED BATH & BEYOND INC	44-45	\$2.5 B
17	LANTRONIX INC	31-33	\$0.3 B
18	COMTECH TELECOMMUN	31-33	\$0.6 B
19	EXTREME NETWORKS INC	31-33	\$2.0 B
20	MOTOROLA SOLUTIONS INC	31-33	\$45.9 B
<i>Panel B: 20 largest firms with zero software engineer payroll share</i>			
1	O'REILLY AUTOMOTIVE INC	44-45	\$47.3 B
2	HERTZ GLOBAL HOLDINGS INC	53	\$11.2 B
3	API GROUP CORP	23	\$5.8 B
4	AMN HEALTHCARE SERVICES INC	56	\$5.8 B
5	SONOCO PRODUCTS CO	31-33	\$5.7 B
6	CELSIUS HOLDINGS INC	31-33	\$5.6 B
7	WINGSTOP INC	72	\$5.2 B
8	NATIONAL BEVERAGE CORP	31-33	\$4.2 B

Continued on next page

Table E.4: Firms with the highest and lowest software engineer payroll shares (continued)

Rank	Firm	NAICS2	Market Cap (2021)
9	ACADEMY SPO AND OU INC	44-45	\$3.9 B
10	AMERICAN STATES WATER CO	22	\$3.8 B
11	FIRSTCASH HOLDINGS INC	52	\$3.6 B
12	INTERPARFUMS INC	31-33	\$3.4 B
13	LATHAM GROUP INC	31-33	\$3.0 B
14	CLEARWAY ENERGY INC	22	\$2.9 B
15	KENON HOLDINGS LTD	22	\$2.8 B
16	H2O AMERICA	22	\$2.2 B
17	KNOWLES CORP	31-33	\$2.2 B
18	XENON PHARMACEUTICALS INC	31-33	\$1.6 B
19	HIGHPEAK ENERGY INC	21	\$1.4 B
20	FLEX LNG LTD	48-49	\$1.3 B

Notes: Panel A reports the 20 firms with the highest software engineer payroll shares in the sample. Panel B reports the 20 largest firms with zero software engineer payroll share in Revelio. Market capitalization is reported in billions of U.S. dollars.

Table E.5: Summary statistics for baseline first-stage firm betas on THNQ excess returns

Statistic	Value
N	1874
Mean	-0.043
SD	0.512
Min	-3.331
P10	-0.529
P25	-0.314
Median	-0.084
P75	0.194
P90	0.541
Max	3.346
Share > 0	0.411

Note: The sample uses the balanced 2022–2025 panel and the baseline FF3 specification. Betas are firm-level coefficients on weekly THNQ excess returns from the baseline FF3 first-stage specification. Each observation is a unique firm (gvkey). *Share > 0* is the fraction of firms with a positive beta.

F Franchise Duration and Heterogeneous SDF Loading

This appendix microfound the extended firm-value response (11) of Section 2.4 by endowing each firm with a finite franchise duration on the common stochastic discount factor.

F.1 Setup

Each firm j is a single legal entity that holds a proprietary franchise over its good. If firm j is alive at the start of period t , it survives to the start of $t + 1$ with probability $\varrho_j \in (0, 1]$, independently across firms and independently of aggregate shocks. Upon death, the incumbent's equity is extinguished and a replacement entity with identical technology (ϕ_j, ϱ_j) is immediately installed in the same goods-market position, facing the same CES demand curve. The replacement's equity is endowed to the representative household at birth. The parameter ϱ_j is a structural primitive of firm j 's franchise (a patent length, imitation lag, or charter renewal horizon), not an asset-pricing parameter. The benchmark of Section 2 corresponds to $\varrho_j = 1$.

For the discount-rate component, write the one-period discount factor as

$$\beta_t = \beta e^{v_t^\Xi}, \quad \Xi_{t,t+1} = \beta e^{v_t^\Xi} \frac{u_c(C_{t+1})}{u_c(C_t)}.$$

The steady state has $v_t^\Xi = 0$, so the steady-state SDF satisfies $\Xi_{t,t+k} = \beta^k$.

Because the replacement firm operates the same technology and faces the same demand curve as the incumbent it succeeds, the cross-sectional distribution of active firms is time-invariant: at each date exactly one firm of each index j is alive and produces using technology ϕ_j . The static allocation, equilibrium factor prices (w_t, W_t) , unit costs mc_{jt} , Domar weights λ_{jt} , and aggregate consumption C_t are identical to those of the benchmark economy. Mortality changes only the mapping from firm j 's profits into equity claims on those profits, not the profits themselves. In particular, the period-profit response of Section 2 carries through verbatim.

The replacement at $t + 1$ has the same technology, faces the same aggregate state, and is itself alive with probability ϱ_j going forward, so its equity at birth has the same conditional value as the incumbent would have had at $t + 1$ had it survived. Firm-level death is therefore idiosyncratic with respect to the household's marginal utility, and the representative-household SDF $\Xi_{t,t+1}$ depends only on aggregate shocks. The same SDF prices every firm's equity. Firm j 's equity, the claim on the operating profits of the current incarnation, satisfies the recursive pricing equation

$$V_{jt} = \pi_{jt} + \varrho_j \mathbb{E}_t[\Xi_{t,t+1} V_{j,t+1}],$$

together with the transversality condition $\lim_{k \rightarrow \infty} \varrho_j^k \mathbb{E}_t[\Xi_{t,t+k} V_{j,t+k}] = 0$, which holds for value paths bounded around the steady state because $\beta \varrho_j < 1$. Iterating forward and using the independence of

survival from aggregate shocks,

$$V_{jt} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \varrho_j^k \Xi_{t,t+k} \pi_{j,t+k} \right]. \quad (\text{E1})$$

Setting $\varrho_j = 1$ recovers (2) exactly.

F.2 Firm-Value Response

Define the firm-specific news operator

$$\mathcal{X}_t^{(j)} \equiv (1 - \tilde{\beta}_j)(\mathbb{E}_t - \mathbb{E}_{t-1}) \left[\sum_{k=0}^{\infty} \tilde{\beta}_j^k d \log X_{t+k} \right], \quad \tilde{\beta}_j \equiv \beta \varrho_j, \quad (\text{E2})$$

for any stochastic process X . When the period- t forecast revision is uniform across horizons, as under a random-walk process for $\log X_t$, $\mathcal{X}_t^{(j)}$ equals the period- t innovation $d \log X_t - \mathbb{E}_{t-1}[d \log X_t]$ for every ϱ_j , because the duration weights $(1 - \tilde{\beta}_j)\tilde{\beta}_j^k$ sum to one. When $\varrho_j = 1$, $\mathcal{X}_t^{(j)}$ reduces to the benchmark news operator $\mathcal{X}_t$.

With mortality, R_{jt} is the return on firm j 's equity conditional on the incumbent surviving the window, which to first order equals $d \log V_{jt} - d \log V_{j,t-1}$, and $\mathbb{E}_{t-1}[R_{jt}]$ is the corresponding survival-conditional expectation. The mortality component of the unconditional return is idiosyncratic and unpriced, and drops out of the aggregate-news revisions below.

Lemma 2 (Firm-Value Response with Franchise Duration). *To first order, firm j 's unexpected return is*

$$R_{jt} - \mathbb{E}_{t-1}[R_{jt}] = \delta^S(\gamma_j - \gamma) \mathcal{S}_t^{(j)} + \mathcal{C}_t^{(j)} + \mathcal{D}_t^{(j)}, \quad (\text{E3})$$

where $\mathcal{S}_t^{(j)}$ and $\mathcal{C}_t^{(j)}$ are the firm-specific news in $\log S$ and $\log C$ defined via (E2), and the firm-specific SDF news is

$$\mathcal{D}_t^{(j)} \equiv (1 - \tilde{\beta}_j)(\mathbb{E}_t - \mathbb{E}_{t-1}) \left[\sum_{k=0}^{\infty} \tilde{\beta}_j^k d \log \Xi_{t,t+k} \right].$$

Proof. The franchise-duration block leaves the static allocation unchanged, so the period-profit response of Section 2 applies horizon by horizon: $d \log \pi_{j,t+k} = \delta^S(\gamma_j - \gamma) d \log S_{t+k} + d \log C_{t+k}$. As elsewhere in the paper, we drop t subscripts when referring to steady-state values. From (E1), the steady-state relations $\Xi_{t,t+k} = \beta^k$ and $\pi_{j,t+k} = \pi_j$ give

$$V_j = \sum_{k=0}^{\infty} (\beta \varrho_j)^k \pi_j = \frac{\pi_j}{1 - \tilde{\beta}_j}.$$

First-order perturbation of (E1) around this steady state yields, for each horizon k ,

$$d(\Xi_{t,t+k} \pi_{j,t+k}) = \beta^k \pi_j (d \log \Xi_{t,t+k} + d \log \pi_{j,t+k}),$$

so

$$dV_{jt} = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \varrho_j^k \beta^k \pi_j (d \log \Xi_{t,t+k} + d \log \pi_{j,t+k}) \right].$$

Dividing by the steady-state value V_j and substituting the period-profit response gives the level response in terms of the present-value operators $(1 - \tilde{\beta}_j)\mathbb{E}_t[\cdot]$. The return satisfies $R_{jt} \approx d \log V_{jt} - d \log V_{j,t-1}$. Subtracting its conditional expectation at $t - 1$ applies the revision $(\mathbb{E}_t - \mathbb{E}_{t-1})$ to each operator and yields (E3). $\square$

Cross-sectional heterogeneity in loadings on aggregate news arises through the effective discount factor $\tilde{\beta}_j$: when an innovation has a non-flat expected horizon profile, longer-duration firms (higher ρ_j) place more weight on distant horizons and load differently than shorter-duration firms.

E.3 Closed-Form Loading

To recover the scalar SDF-news representation used in the main text, we impose two conditions.

Assumption 3. (*Uniform real cash-flow news*) At date t , the forecast revisions are uniform across horizons: $(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log S_{t+k}] = \mathcal{S}_t$ and $(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log C_{t+k}] = \mathcal{C}_t$ for all $k \geq 0$, where $\mathcal{S}_t$ and $\mathcal{C}_t$ are the benchmark ($\rho_j = 1$) news in $\log S$ and $\log C$.

For the SDF news, decompose $d \log \Xi_{t,t+k} = d \log \Xi_{t,t+k}|\Xi + d \log u_c(C_{t+k}) - d \log u_c(C_t)$, where $d \log \Xi_{t,t+k}|\Xi \equiv \sum_{\ell=0}^{k-1} dv_{t+\ell}^\Xi$ is the discount-rate component.

Assumption 4. (*One-factor SDF term structure*) Let $\epsilon_t^\Xi \equiv dv_t^\Xi - \mathbb{E}_{t-1}[dv_t^\Xi]$ denote the period- t innovation in the discount-factor state. The revision in the expected cumulative log-SDF response to this innovation admits the representation

$$(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log \Xi_{t,t+k}|\Xi] = a_k \epsilon_t^\Xi, \quad a_0 = 0,$$

for a deterministic sequence $\{a_k\}_{k \geq 0}$ common across firms, with $\mathbb{E}_{t-1}[\epsilon_t^\Xi] = 0$, $\sum_{k=0}^{\infty} \beta^k |a_k| < \infty$, and $(1 - \beta) \sum_{k=0}^{\infty} \beta^k a_k \neq 0$.

Assumption 3 means that period- t news shifts the expected paths of $\log S$ and $\log C$ uniformly at all horizons. Because a uniform revision passes through the unit-sum weights in (E2) unchanged, its common value automatically equals the benchmark news $\mathcal{S}_t$ (resp. $\mathcal{C}_t$). The assumption is the appendix counterpart of the random-walk normalization behind (6): under a random-walk process for $\log S_t$, the period- t innovation $d \log S_t - \mathbb{E}_{t-1}[d \log S_t]$ revises expected productivity by the same amount at every horizon, so the assumption holds with $\mathcal{S}_t$ equal to that innovation. Dynamics already anticipated at $t - 1$, such as a predictable drift, are left unrestricted. Under Assumption 3, the consumption component of SDF news drops out of the period- t revision to first order, leaving only the discount-rate component. Assumption 4 then requires the horizon profile of this remaining news to be generated by the single innovation ϵ_t^Ξ , with common term-structure loadings $\{a_k\}$. This condition holds, for example, when v_t^Ξ follows a linear stationary process driven by a single innovation. Define

the benchmark discount-rate news by applying the duration weights at $\varrho_j = 1$ to the discount-rate component of the SDF:

$$\mathcal{D}_t \equiv (1 - \beta)(\mathbb{E}_t - \mathbb{E}_{t-1}) \left[\sum_{k=0}^{\infty} \beta^k d \log \Xi_{t,t+k} | \Xi \right]. \quad (\text{E4})$$

Proposition 7 (Heterogeneous SDF Loading). *Under Assumptions 3 and 4, the unexpected firm return admits the representation*

$$R_{jt} - \mathbb{E}_{t-1}[R_{jt}] = \delta^S \gamma_j \mathcal{S}_t + \eta_j \mathcal{D}_t + \mu_t, \quad (\text{E5})$$

where the firm-invariant component is $\mu_t \equiv \mathcal{C}_t - \delta^S \gamma \mathcal{S}_t$, and the loading is

$$\eta_j = \frac{(1 - \tilde{\beta}_j) \sum_{k=0}^{\infty} \tilde{\beta}_j^k a_k}{(1 - \beta) \sum_{k=0}^{\infty} \beta^k a_k}, \quad \tilde{\beta}_j = \beta \varrho_j, \quad (\text{E6})$$

depending only on ϱ_j and the common term structure $\{a_k\}$. When the discount-factor state follows the $AR(1)$ $dv_t^\Xi = \rho_\Xi dv_{t-1}^\Xi + \epsilon_t^\Xi$, with persistence $\rho_\Xi \in [0, 1)$,

$$\eta_j = \frac{\varrho_j(1 - \beta \rho_\Xi)}{1 - \beta \varrho_j \rho_\Xi}, \quad (\text{E7})$$

which is strictly increasing in ϱ_j . The benchmark $\varrho_j = 1$ gives $\eta_j = 1$.

Proof. Under Assumption 3, the revision inside (E2) equals $\mathcal{S}_t$ (resp. $\mathcal{C}_t$) at every horizon, and the duration weights $(1 - \tilde{\beta}_j) \tilde{\beta}_j^k$ sum to one, so $\mathcal{S}_t^{(j)} = \mathcal{S}_t$ and $\mathcal{C}_t^{(j)} = \mathcal{C}_t$ independent of j . Substituting into (E3) and collecting firm-invariant terms gives

$$R_{jt} - \mathbb{E}_{t-1}[R_{jt}] = \delta^S \gamma_j \mathcal{S}_t + \mathcal{D}_t^{(j)} + \underbrace{(\mathcal{C}_t - \delta^S \gamma \mathcal{S}_t)}_{=\mu_t}.$$

By Assumption 3, the consumption component of the SDF news drops out of the revision: to first order, $d \log u_c(C_{t+k}) - d \log u_c(C_t) = \frac{u_{cc}(C)C}{u_c(C)} (d \log C_{t+k} - d \log C_t)$, where C is steady-state consumption, and the revisions of $d \log C_{t+k}$ and $d \log C_t$ both equal $\mathcal{C}_t$, so $(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log C_{t+k} - d \log C_t] = 0$ for all $k \geq 0$. Hence $(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log \Xi_{t,t+k}] = (\mathbb{E}_t - \mathbb{E}_{t-1})[d \log \Xi_{t,t+k} | \Xi]$. Assumption 4 then gives $\mathcal{D}_t^{(j)} = (1 - \tilde{\beta}_j) \sum_{k=0}^{\infty} \tilde{\beta}_j^k a_k \cdot \epsilon_t^\Xi$, and (E4) gives $\mathcal{D}_t = (1 - \beta) \sum_{k=0}^{\infty} \beta^k a_k \cdot \epsilon_t^\Xi$. Both are proportional to ϵ_t^Ξ ; defining η_j as the ratio of their coefficients gives (E6), which depends on ϱ_j but not on ϵ_t^Ξ .

For the $AR(1)$ specialization, the conditional forecast revision satisfies

$$(\mathbb{E}_t - \mathbb{E}_{t-1})[dv_{t+\ell}^\Xi] = \rho_\Xi^\ell \epsilon_t^\Xi.$$

Therefore,

$$(\mathbb{E}_t - \mathbb{E}_{t-1})[d \log \Xi_{t,t+k} | \Xi] = \sum_{\ell=0}^{k-1} \rho_\Xi^\ell \epsilon_t^\Xi = \frac{1 - \rho_\Xi^k}{1 - \rho_\Xi} \epsilon_t^\Xi,$$

which is Assumption 4 with $a_k = (1 - \rho_\Xi^k)/(1 - \rho_\Xi)$. Evaluating the numerator of (E6),

$$(1 - \tilde{\beta}_j) \sum_{k=0}^{\infty} \tilde{\beta}_j^k a_k = \frac{1 - \tilde{\beta}_j}{1 - \rho_\Xi} \left(\frac{1}{1 - \tilde{\beta}_j} - \frac{1}{1 - \tilde{\beta}_j \rho_\Xi} \right) = \frac{\tilde{\beta}_j}{1 - \tilde{\beta}_j \rho_\Xi}.$$

The denominator evaluates analogously with $\tilde{\beta}_j$ replaced by β , giving $\beta/(1 - \beta\rho_\Xi)$. Taking the ratio and substituting $\tilde{\beta}_j = \beta\varrho_j$,

$$\eta_j = \frac{\tilde{\beta}_j/(1 - \tilde{\beta}_j\rho_\Xi)}{\beta/(1 - \beta\rho_\Xi)} = \frac{\varrho_j(1 - \beta\rho_\Xi)}{1 - \beta\varrho_j\rho_\Xi},$$

which is (E7). Differentiating with respect to ϱ_j ,

$$\frac{\partial \eta_j}{\partial \varrho_j} = \frac{1 - \beta\rho_\Xi}{(1 - \beta\varrho_j\rho_\Xi)^2} > 0$$

for all $\beta \in (0, 1)$, $\rho_\Xi \in [0, 1)$, and $\varrho_j \in (0, 1]$. □

Under the AR(1) specialization, the loading η_j is monotone in ϱ_j : longer-duration firms load more strongly on common discount-rate news. Two limiting cases are illuminating. With $\rho_\Xi = 0$ (iid news), $\eta_j = \varrho_j$, so the loading equals the survival probability directly. With $\rho_\Xi \rightarrow 1$ (permanent news), $\eta_j = \varrho_j(1 - \beta)/(1 - \beta\varrho_j)$, the familiar duration-of-a-perpetuity formula adjusted for franchise mortality.

G Production Network Economy

This appendix sets up the production network economy of Section 4.3, states the results used there, and documents the measurement and calibration behind the reported numbers. The economy places a continuum of firms inside each of N sectors: firms are the objects the empirical analysis observes, sectors are the objects the input-output accounts measure, and an exact aggregation result connects the two tiers. All results established in Section 2 for a single sector are taken as given. Proofs are provided at the end of this appendix.

G.1 Environment

The economy comprises N sectors, indexed by n , each populated by a unit continuum of firms indexed by $j \in [0, 1]$. The household, the inelastically supplied endowments s and L , the complete-markets structure, and the timing are as in Section 2. The economy is driven by software engineering productivity S_t and by a vector Z_t of other aggregate shocks with components Z_{kt} indexed by k , which may include Hicks-neutral productivity, discount-factor shocks, and shifters of the composition of demand across sectors, and which may be arbitrarily correlated.

Demand. Consumption is a two-tier CES aggregate. The upper tier combines sector composites C_{nt} with elasticity ζ and demand weights $\omega_n(Z_t)$ summing to one; the lower tier combines the varieties of the firms in a sector with the elasticity ε of Section 2,

$$C_t = \left[\sum_{n=1}^N \omega_n(Z_t)^{1/\zeta} C_{nt}^{(\zeta-1)/\zeta} \right]^{\zeta/(\zeta-1)}, \quad C_{nt} = \left[\int_0^1 c_{nt}(j)^{(\varepsilon-1)/\varepsilon} dj \right]^{\varepsilon/(\varepsilon-1)},$$

with sector price indices P_{nt} , consumer price index $P_t = 1$ as numeraire, and final-expenditure shares $f_{nt} = P_{nt}C_{nt}/\sum_m P_{mt}C_{mt}$ summing to one. When $\zeta = \varepsilon$, the two CES demand tiers collapse to the single CES demand aggregator used in Section 2.

The sector composite is the only channel through which sector n 's output reaches the economy: the household consumes the quantity C_{nt} of it, and every downstream firm purchases the same bundle as an intermediate input, allocating across the constituent varieties with the same elasticity ε . This is the substantive aggregation assumption of the model, and it is what the nesting of Section 2 already implies: a firm that cuts its relative price gains demand from its sector peers at the rate ε whether the buyer is a household or a downstream producer.

Production. Firm j in sector n produces with the technology of Section 2, with sector composites in place of final-good materials,

$$y_{nt}(j) = A_t z_n(j) \left(\frac{H_{nt}(j)}{\alpha_n} \right)^{\alpha_n} \prod_{m: \Omega_{nm} > 0} \left(\frac{X_{nm,t}(j)}{\Omega_{nm}} \right)^{\Omega_{nm}}, \quad \alpha_n + \sum_{m=1}^N \Omega_{nm} = 1,$$

where $X_{nm,t}(j)$ is the firm's use of the sector- m composite, Ω_{nm} the corresponding cost share, and $H_{nt}(j)$ the CES labor composite of Section 2, with software engineering intensity $\phi_n(j)$, substitution elasticity σ , and unit cost $q_{nt}(j)$, so that marginal cost is $mc_{nt}(j) = q_{nt}(j)^{\alpha_n} \prod_m P_{mt}^{\Omega_{nm}} / (A_t z_n(j))$. The intermediate block extends Section 2's final-good materials to the multi-sector setting. Capital is a reproducible input rather than a separate accumulated factor: firms purchase capital services through the same intermediate block, so the row Ω_n absorbs materials and capital services alike. Accordingly, α_n is the labor share of sector n 's total costs, and GDP is value added. Firms in a sector share $(\alpha_n, \Omega_n, \sigma)$ and differ only in $\phi_n(j)$, which generates the payroll share $\gamma_{nt}(j) = w_t s_{nt}(j) / (w_t s_{nt}(j) + W_t L_{nt}(j))$, and in the idiosyncratic productivity $z_n(j)$, which is orthogonal to $\phi_n(j)$.

Each firm is atomistic within its sector: it takes sector composites and aggregate prices as given, faces the within-sector CES demand with elasticity ε , and charges the constant markup $\mu = \varepsilon / (\varepsilon - 1)$ of Section 2 (a parameter, not to be confused with the firm-invariant return component μ_t), so $p_{nt}(j) = \mu \cdot mc_{nt}(j)$ and $\pi_{nt}(j) = \frac{1}{\varepsilon} p_{nt}(j) y_{nt}(j)$.

Sector aggregates. Let $\mathcal{R}_{nt}(j) = p_{nt}(j) y_{nt}(j)$ denote firm sales, $\mathcal{R}_{nt} = \int_0^1 \mathcal{R}_{nt}(j) dj$ aggregate sector sales, and $\varsigma_{nt}(j) = \mathcal{R}_{nt}(j) / \mathcal{R}_{nt} = (p_{nt}(j) / P_{nt})^{1-\varepsilon}$ the firm's share of sector sales. We summarize the within-sector distribution of payroll shares by its sales-weighted mean and variance,

$$\bar{\gamma}_{nt} = \int_0^1 \varsigma_{nt}(j) \gamma_{nt}(j) dj, \quad \sigma_{\gamma,n}^2 = \int_0^1 \varsigma_{nt}(j) (\gamma_{nt}(j) - \bar{\gamma}_{nt})^2 dj,$$

the sector-level counterparts of γ and σ_γ^2 in Section 2. Sector n 's software engineering and non-SWE labor cost shares are $\alpha_n \bar{\gamma}_n$ and $\bar{\theta}_n = \alpha_n (1 - \bar{\gamma}_n)$. Stacking the intermediate shares into the $N \times N$ matrix Ω , the Leontief inverse is $\Lambda = (\mathbf{I} - \Omega)^{-1}$, the cost-share identity gives $\Lambda \alpha = \mathbf{1}$, and the cost-based Domar

weight $\tilde{\lambda} = \Lambda^\top f$ propagates final expenditure through the un-marked-up Leontief inverse, coinciding with the sales share $\mathcal{R}_n/(P_t C_t)$ when $\mu = 1$ or $\Omega = 0$. Finally, we record the equilibrium factor-price responses to a unit software engineering productivity shock in the two scalars

$$u = \frac{d \log w_t - d \log S_t}{d \log S_t}, \quad v = \frac{d \log W_t}{d \log S_t},$$

so that $u - v$ is the response of the effective relative wage that drives the entire cross-section.

G.2 Aggregation, Identification, and the GDP Impact

The first result shows that the sector-level objects measured in the input-output accounts are exactly the sales-weighted aggregates of their firm-level counterparts and do not otherwise depend on the shape of the within-sector distribution.

Lemma 3 (Exact Aggregation). *Suppose firms in sector n share α_n and the common markup μ . Then:*

- (i) *The sector software engineering cost share, sector SWE compensation divided by sector total cost, equals the sales-weighted mean of the firm cost shares, $\int_0^1 \zeta_n(j) \alpha_n \gamma_n(j) dj = \alpha_n \bar{\gamma}_n$.*
- (ii) *Each firm's cost-based Domar weight factorizes as $\tilde{\lambda}_n(j) = \tilde{\lambda}_n \zeta_n(j)$, so the economy-wide cost-based-Domar-weighted SWE cost share aggregates exactly to the sector level,*

$$\sum_{n=1}^N \int_0^1 \tilde{\lambda}_n(j) \alpha_n \gamma_n(j) dj = \sum_{n=1}^N \tilde{\lambda}_n \alpha_n \bar{\gamma}_n.$$

The sales weighting is a consequence of the common composite: a firm's weight in any sector aggregate is its share of sector sales, because that is the share of every transaction the sector supplies. Since a sales-weighted mean of ratios collapses to a ratio of totals, the sector share in part (i) is the sector SWE wage bill over sector cost, with the correct weighting already built in. This aggregation result allows us to measure sector levels and firm slopes using different data sources.

We next characterize the first-order response of firm profits around the steady state, dropping t subscripts from steady-state values. Let $v_n = \alpha_n \bar{\gamma}_n \mathcal{R}_n / \sum_k \alpha_k \bar{\gamma}_k \mathcal{R}_k$ denote sector n 's share of economy-wide SWE compensation and $v_n^L = \bar{\theta}_n \mathcal{R}_n / \sum_k \bar{\theta}_k \mathcal{R}_k$ its share of non-SWE compensation.

Lemma 4 (Equilibrium Responses). *To first order in $d \log S_t$, under the normalization $d \log S_t = 1$:*

- (i) *A firm's price relative to its sector composite moves only with its own software engineering intensity, $d \log p_{nt}(j) - d \log P_{nt} = \alpha_n (\gamma_n(j) - \bar{\gamma}_n) (u - v)$.*
- (ii) *Net of the economy-wide expenditure response, the firm's profit response is a sector-common term plus a firm-specific deviation carrying the within-sector demand elasticity,*

$$d \log \pi_{nt}(j) - d \log E_t = b_n + (1 - \varepsilon) \alpha_n (\gamma_n(j) - \bar{\gamma}_n) (u - v), \quad (\text{G.1})$$

where $b_n = d \log \mathcal{R}_n - d \log E_t$ is common to all firms in sector n and carries the input-output and across-sector propagation relative to aggregate expenditure: stacking, $b = \mathcal{B}(\alpha \bar{\gamma} u + \bar{\theta} v)$, where the sector sales-response matrix $\mathcal{B}$ is built from the Leontief inverse Λ , the across-sector elasticity ζ , and the demand-side inverse $\Psi = (\mathbf{I} - \mathcal{T})^{-1}$, with $\mathcal{T}_{nm} = \Omega_{mn} \mathcal{R}_m / (\mu \mathcal{R}_n)$ the share of sector n 's sales accounted for by sector m 's intermediate purchases.

- (iii) The two factor-market-clearing conditions and the numeraire condition jointly determine $(u, v, d \log E_t)$. Eliminating $d \log E_t$ leaves a 2×2 linear system in (u, v) , in which the within-sector distributions enter only through the moments $(\bar{\gamma}_n, \sigma_{\gamma, n}^2)$.

Part (i) is the key for the aggregation: all firms in a sector buy the same composites at the same shares, so their marginal costs differ only through the labor composite, and within it only through the payroll share. Part (ii) turns the price deviation into a profit deviation through the within-sector elasticity, since ε is the rate at which every buyer substitutes inside the composite; the across-sector elasticity ζ and the network matrices (Ψ, Λ) appear only in b_n , which the second stage's sector fixed effects absorb.

Capitalizing the profit responses from Lemma 4 into firm values, as in Section 2, yields firm (n, j) 's unexpected period- t return:

$$R_{nt}(j) - \mathbb{E}_{t-1}[R_{nt}(j)] = b_n \mathcal{S}_t + \sum_k b_n^{Z_k} \mathcal{Z}_{kt} + (1 - \varepsilon) \alpha_n (\gamma_n(j) - \bar{\gamma}_n) \left[(u - v) \mathcal{S}_t + \sum_k (u_k - v_k) \mathcal{Z}_{kt} \right] + \mu_t, \quad (\text{G.2})$$

where $\mathcal{Z}_{kt}$ is the news in normalized present-value Z_k defined analogously to $\mathcal{S}_t$, the scalars (u_k, v_k) are the factor-price responses to a unit Z_k shock, $b_n^{Z_k}$ collects the sector-common response to Z_k net of common expenditure, and μ_t collects firm-invariant news, including revisions to capitalized common expenditure. Every within-sector deviation carries the same gradient, because inside a sector an aggregate shock can separate firms by software engineering intensity only by moving the effective relative wage. Under condition (ii) in Section 2.3, unexpected returns and excess returns have the same first-stage AI-index beta. Substituting (G.2) into the two-stage regression and projecting on the payroll share conditional on the controls X_j , which include the sector fixed effects, yields

$$\delta = \kappa \delta_{\text{in}}^S + \sum_k \kappa_k^Z \delta_k^Z, \quad \delta_{\text{in}}^S = (1 - \varepsilon) \bar{\alpha} (u - v), \quad \delta_k^Z = (1 - \varepsilon) \bar{\alpha} (u_k - v_k), \quad (\text{G.3})$$

where κ and κ_k^Z are the projection coefficients of $(\mathcal{S}_t, \mathcal{Z}_{kt})$ onto the residualized AI-index return. Here $\bar{\alpha}$ is the coefficient on the SWE payroll share γ_j in a regression of the SWE cost share $\alpha_{n(j)} \gamma_j$ on γ_j , using the same sample, controls, and fixed effects as the empirical second stage. It measures how differences in payroll shares translate into differences in cost shares after applying those controls. The sector fixed effects absorb b_n and $b_n^{Z_k}$, so network propagation does not affect the estimated slope. Because the response is affine in $\gamma_n(j)$ within each model sector, using NAICS3 fixed effects

leaves the within-cell slope unchanged when NAICS3 cells are at least as fine as the model's sectors. Identification then requires only that the other aggregate shocks not move the effective relative wage.

Assumption 5. (*Balanced incidence in the network.*) For each aggregate shock Z_k , the sector sales responses are balanced across software-engineering-intensive and other sectors,

$$\sum_{n=1}^N (v_n - v_n^L) b_n^{Z_k} = 0.$$

Assumption 5 is the network analogue of the second condition in Assumption 1. The first condition is unnecessary here because the sector fixed effects in X absorb each shock's direct sector-level incidence. It holds exactly for any shock whose sector-level incidence is uniform, such as discount-factor shocks and, under a common labor share, Hicks-neutral productivity. Its main content is to restrict demand-composition shocks: expenditure shifts across sectors must not systematically favor sectors that employ software engineers intensively. Unlike Assumption 1, which restricts primitives, it is stated on the equilibrium incidence $b_n^{Z_k}$, because that is exactly what identification requires.

Proposition 8 (Identification in the Network Economy). Under Assumption 5, $u_k = v_k$ for every k , so $\delta_k^Z = 0$, the equation (G.3) collapses to $\delta = \kappa \delta_{in}^S$, and

$$\kappa = \frac{\delta}{\delta_{in}^S}, \quad \delta_{in}^S = (1 - \varepsilon) \bar{\alpha} (u - v). \quad (\text{G.4})$$

In the one-sector roundabout economy of Section 2 ($N = 1$ and $\Omega_{11} = 1 - \alpha$, so that $b = 0$), the system of Lemma 4(iii) gives $u - v = -1/\Sigma$ and $\delta_{in}^S = \alpha(\varepsilon - 1)/\Sigma$, the closed form of Section 2.

The two-stage procedure therefore recovers κ by dividing the empirical slope by the corresponding model-implied slope. In the network economy, this model-implied slope depends on the equilibrium relative-wage response $u - v$ and the coefficient $\bar{\alpha}$ defined above.

Turning to the GDP impact, real GDP is the consumption aggregate $Y_t = C_t$. Define the cost-based-Domar-weighted primary-factor technology weights $\tilde{s}_S = \sum_n \tilde{\lambda}_n \alpha_n \tilde{\gamma}_n$ and $\tilde{s}_L = \sum_n \tilde{\lambda}_n \bar{\theta}_n$. They sum to one because $\tilde{s}_S + \tilde{s}_L = f^\top \Lambda \alpha = 1$. Let $s_{S,t} = w_t s / Y_t$ and $s_{L,t} = W_t L / Y_t$ denote the actual factor-income shares in GDP.

Proposition 9 (Modified Hulten Formula in the Network Economy). To first order, the GDP impact of the software engineering channel is

$$d \log Y_t|_S = \gamma^{IO} d \log S_t, \quad \gamma^{IO} = \underbrace{\sum_{n=1}^N \tilde{\lambda}_n \alpha_n \tilde{\gamma}_n}_{\text{cost-based Domar technology}} - \underbrace{\sum_{\tau \in \{S, L\}} \tilde{s}_\tau \frac{d \log s_{\tau, t}}{d \log S_t}}_{\text{factor reallocation } \Delta}, \quad (\text{G.5})$$

a cost-based-Domar-weighted sum of sector software engineering cost shares, minus the Baqaee and Farhi (2020) factor-reallocation term. The reallocation term vanishes when $\mu = 1$ or $\Omega = 0$. Aggregating over horizons as in Section 2 gives $\mathcal{Y}_t = \gamma^{IO} \mathcal{S}_t$. The best linear predictor of GDP news given the

residualized AI-index return is

$$\mathbb{E}^*[\mathcal{Y}_t | \tilde{R}_t^{AI}] = \gamma^{IO} \kappa \cdot \tilde{R}_t^{AI}. \quad (\text{G.6})$$

G.3 Taking the Model to the Data

The calculation in the next subsection requires the production network and final-demand weights, sector labor and software engineering shares, the elasticities and markup, and two empirical inputs: the second-stage slope and the cumulative residualized AI-index return. We calibrate each input as follows.

The production network and labor shares. We use the 2024 BEA Summary input-output tables and retain $N = 65$ private sectors. We exclude owner-occupied housing and five government sectors, so the resulting GDP effect refers to the private market economy. The accounts supply the observed intermediate transactions and sector gross output, from which we construct the input cost-share matrix Ω and residual final-demand shares f . We obtain the sector labor cost share from

$$\alpha_n = \mu \times \frac{\text{labor compensation}_n}{\text{gross output}_n}.$$

The gross-up converts the revenue share in the accounts into the cost share used by the production function. We treat capital services and imports as produced inputs that can expand with the economy, rather than as fixed primary factors. Because they do not appear as separate supplying sectors in our private-sector network, we allocate them across each buyer's observed suppliers in the same proportions as its other intermediate purchases. We therefore rescale each row of Ω to sum to $1 - \alpha_n$, which ensures $\Lambda\alpha = \mathbf{1}$. After this adjustment, the original BEA sector-output figures no longer exactly balance the model's supplies and demands. We resolve this discrepancy by holding the adjusted production network and the observed composition of final demand fixed and solving for the sector outputs that clear all goods markets:

$$\mathcal{R} = (\mathbf{I} - \Omega^\top / \mu)^{-1} f E,$$

where the scale E is immaterial. We also recompute each sector's final-sales share from this equilibrium. Sector output in the calibrated model may therefore differ slightly from its BEA counterpart. Both the linearized calculation and the nonlinear calculation use this same recalibrated economy, so their results are directly comparable.

Elasticities and markup. We set the within-firm elasticity between software engineering and other labor to $\sigma = 1.6$, as in Section 3.4. The within-sector demand elasticity is $\varepsilon = 5$, implying the uniform markup $\mu = \varepsilon / (\varepsilon - 1) = 1.25$. We set the across-sector elasticity to $\zeta = \varepsilon = 5$, so the two CES demand tiers collapse to a single CES aggregator. We impose Cobb-Douglas intermediate-input shares.

Sector software engineering shares. We map the all-sector Compustat-Revelio firm panel into the BEA sectors and use 2021 sales as weights. We retain firms with a nonmissing payroll share that map to a retained sector; firms with zero or missing 2021 sales receive zero weight. For each sector n , we compute

$$\bar{\gamma}_n = \frac{\sum_j \mathcal{R}_n(j) \gamma_n(j)}{\sum_j \mathcal{R}_n(j)}, \quad \sigma_{\gamma,n}^2 = \frac{\sum_j \mathcal{R}_n(j) (\gamma_n(j) - \bar{\gamma}_n)^2}{\sum_j \mathcal{R}_n(j)}.$$

These moments enter the factor-market-clearing system, while $\alpha_n \bar{\gamma}_n$ enters the GDP weight. We use the all-sector panel because the GDP calculation must include software-producing sectors that are excluded from the regression sample. Eight sectors with no assigned firms are given $\bar{\gamma}_n = \sigma_{\gamma,n}^2 = 0$. For the nonlinear calculation, we assume that $\gamma_n(j)$ follows a Beta distribution, a standard distribution for shares bounded between zero and one, whose two parameters are uniquely chosen to match $(\bar{\gamma}_n, \sigma_{\gamma,n}^2)$; a zero-variance sector has a degenerate distribution at its mean. This assumption introduces no additional calibrated parameter beyond the two moments used in the linearized calculation.

The empirical slope and AI returns. The baseline second-stage estimate is $\delta = 1.24$, from column 4 of Table 1, which includes NAICS3, asset-size-quintile, and sales-size-quintile fixed effects. To obtain the structural gradient δ_{in}^S , we run the same regression inside the model on firms in the regression sample that can be uniquely assigned to a retained BEA Summary sector, using the same payroll-share regressor, fixed effects, and controls. Finally, following Section 3.4, the cumulative residualized AI-index return is $\sum_{t=\text{Nov. 2022}}^{\text{Dec. 2025}} \tilde{R}_t^{\text{AI}} = 22.9\%$.

G.4 Numerical Procedure and Results

We now use the calibrated production network model to infer software engineering productivity news from the AI-index return and quantify the implied effect on GDP. The calculation proceeds in three steps. Factor prices are obtained by combining the sector price map through Λ with the demand map $\mathcal{B}$ of Lemma 4(ii). Substituting these maps into the factor-market-clearing conditions and imposing the numeraire eliminates the common expenditure response, leaving a 2×2 linear system for (u, v) . The structural gradient δ_{in}^S is obtained by computing each included firm's profit elasticity from (G.1) and running the same payroll-share regression, with the same controls and fixed effects, as in the empirical second stage. The GDP calculation uses the cost-based Domar weights $\tilde{\lambda} = \Lambda^\top f$. The technology term is $\tilde{\lambda}^\top (\alpha \tilde{\gamma})$, while the income-side equilibrium response gives the GDP weight γ^{IO} ; their difference is the factor-reallocation term Δ in (G.5). Every object in this linearized system is evaluated at the same internally consistent economy described above.

At the baseline calibration $(\sigma, \varepsilon, \zeta, \mu) = (1.6, 5, 5, 1.25)$, the factor-market-clearing system gives $u -$

$\nu = -0.573$, and the inversion and the GDP impact follow from (G.4) and (G.5):

$$\begin{aligned}\delta_{\text{in}}^S &= (1 - \varepsilon)\bar{\alpha}(u - \nu) = (1 - 5) \times 0.424 \times (-0.573) = 0.97, \\ \kappa &= \delta / \delta_{\text{in}}^S = 1.24 / 0.97 = 1.28, \\ \sum_t \mathbb{E}^*[\mathcal{S}_t | \tilde{R}_t^{\text{AI}}] &= \kappa \sum_t \tilde{R}_t^{\text{AI}} = 1.28 \times 22.9\% \approx 29.3\%, \\ \gamma^{\text{IO}} &= \tilde{\lambda}^\top (\alpha \bar{\gamma}) - \Delta = 0.0850 - 0.0025 = 0.0825, \\ \sum_{t=\text{Nov. 2022}}^{\text{Dec. 2025}} \mathbb{E}^*[\mathcal{Y}_t | \tilde{R}_t^{\text{AI}}] &= \gamma^{\text{IO}} \kappa \sum_t \tilde{R}_t^{\text{AI}} = 0.0825 \times 1.277 \times 22.9\% \approx 2.41\%.\end{aligned}$$

Relative to the one-sector benchmark, the network estimate lowers implied software engineering productivity news from 32.6% to 29.3% and the GDP impact from 3.61% to 2.41%. Table G.1 decomposes this difference. Rows 1–4 start from the one-sector benchmark and sequentially introduce the components of the network calculation. Row 1 reproduces the one-sector benchmark from Section 4.1. Row 2 replaces the benchmark structural gradient $\delta^S = 0.87$ with the network-economy in-sample gradient $\delta_{\text{in}}^S = 0.97$, while holding the empirical slope $\delta = 1.24$ and the benchmark GDP weight 0.1110 fixed. Row 3 retains the Row 2 inversion for software engineering productivity news and replaces the benchmark technology term 0.1110 with its network-economy counterpart, the cost-based-Domar technology term in (G.5), 0.0850. Row 4 subtracts the markup-induced factor-reallocation term in (G.5), $\Delta = 0.0025$, yielding the network-economy GDP elasticity $\gamma^{\text{IO}} = 0.0825$.

The network-economy adjustment to implied software engineering productivity news is modest. Replacing the one-sector structural gradient 0.87 with the network-economy gradient 0.97 lowers implied software engineering productivity news from 32.6% to 29.3% and the GDP impact by 0.37 percentage points, to 3.25%. Aggregation accounts for most of the remaining difference. Replacing the benchmark economy’s cost-based GDP weight of 0.1110 with the network-economy cost-based-Domar technology term in (G.5) lowers the GDP impact by another 0.76 percentage points, to 2.49%. Subtracting the markup-induced factor-reallocation term $\Delta = 0.0025$ lowers the GDP impact by a further 0.07 percentage points. Thus, the inversion accounts for about 30% of the gap between the one-sector and network GDP-impact estimates, while aggregation to GDP accounts for about 70%.

Sensitivity and verification. As in Section 3.4, the results are most sensitive to the within-firm substitution elasticity σ . Varying σ from 0.5 to 2 changes the implied GDP impact in the network economy from 1.06% to 2.91% (Table G.2). The baseline impact is nearly invariant to the across-sector elasticity ζ : lowering ζ from 5 to 0.5 changes the GDP impact only from 2.41% to 2.35% (Table G.3, Panel A). The within-sector elasticity ε matters more for the inversion. Varying ε from 4 to 6 lowers the implied productivity change from 36.3% to 24.6% and the GDP impact from 2.96% to 2.05% (Panel B). The GDP weight itself remains close to 0.082, so this sensitivity operates primarily through the structural gradient and the inferred productivity change. The GDP impact rises from 2.41% to 2.96% if α_n is

Table G.1: From the one-sector benchmark to the production-network estimates

Calculation	Productivity inversion			Aggregation to GDP	
	Empirical slope δ	Structural gradient δ^S	Productivity news	GDP weight	GDP impact
<i>Panel A. Sequential decomposition</i>					
1. One-sector benchmark	1.24	0.87	32.6%	0.1110	3.61%
2. Network-economy structural gradient	1.24	0.97	29.3%	0.1110	3.25%
3. Economy-wide cost-based Domar weight	1.24	0.97	29.3%	0.0850	2.49%
4. Markup-induced factor reallocation	1.24	0.97	29.3%	0.0825	2.41%

Notes: In each row, productivity news is $\sum_t \mathbb{E}^*[\mathcal{S}_t | \tilde{R}_t^{\text{AI}}] = (\delta/\delta^S) \sum_t \tilde{R}_t^{\text{AI}}$, using the entries in the empirical-slope and structural-gradient columns. The GDP impact equals productivity news times the entry in the GDP-weight column. The corresponding one-sector and network-economy weights are denoted by γ and γ^{IO} , respectively. All sums run over the event window from November 2022 through December 2025, and $\sum_t \tilde{R}_t^{\text{AI}} = 22.9\%$. Panel A holds the baseline empirical slope $\delta = 1.24$ fixed and sequentially introduces the components of the network calculation. Row 1 reproduces the one-sector benchmark of Section 4.1. Row 2 replaces only the benchmark structural gradient with the network-economy gradient $\delta_{\text{in}}^S = 0.97$. Row 3 replaces the benchmark GDP weight 0.1110 with the network-economy cost-based-Domar technology term 0.0850 in (G.5). Row 4 subtracts the markup-induced factor-reallocation term $\Delta = 0.0025$, yielding the network-economy GDP elasticity $\gamma^{\text{IO}} = 0.0825$.

measured as the labor share of sales rather than the labor share of total cost.

We verify the first-order solution derived above using a nonlinear general-equilibrium solver for the same calibrated economy. Within each sector, the solver represents firms using the Beta distribution that matches the Revelio sales-weighted mean and variance. It solves the model in levels, rebates profits to the household, and imposes all market-clearing conditions. At the inferred finite productivity change, the exact solution raises the GDP impact only from 2.41% to 2.55% (Table G.4).

Table G.2: GDP impact sensitivity to the within-firm substitution elasticity (production network economy)

σ	δ_{in}^S	κ	Implied productivity change (%)	Implied GDP change (%)
0.5	2.13	0.58	13.4	1.06
1.0	1.38	0.90	20.6	1.68
1.5	1.02	1.21	27.8	2.29
1.6 (baseline)	0.97	1.28	29.3	2.41
2.0	0.81	1.53	35.1	2.91

Notes: Each row solves the production network economy of Appendix G at the indicated σ , holding $\varepsilon = \zeta = 5$, $\mu = \varepsilon/(\varepsilon - 1) = 1.25$, and the baseline slope $\delta = 1.24$ (payroll-share basis). δ_{in}^S is the in-sample structural gradient from Proposition 8 and $\kappa = \delta/\delta_{\text{in}}^S$ the software engineering productivity news per unit of AI-index return. The GDP news through the software engineering channel is $\gamma^{\text{IO}} \kappa \cdot \sum_{t='22 \text{ Nov}}^{\text{end '25}} \tilde{R}_t^{\text{AI}}$, using $\sum_{t='22 \text{ Nov}}^{\text{end '25}} \tilde{R}_t^{\text{AI}} = 22.9\%$ and the weight γ^{IO} of Proposition 9 recomputed at each σ .

Table G.3: Sensitivity to demand elasticities in the production network economy

Parameter	Value	δ_{in}^S	Implied productivity change (%)	GDP weight	Implied GDP change (%)
<i>Panel A. Across-sector demand elasticity</i>					
ζ	0.5	1.03	27.5	0.085	2.35
	1.0	1.03	27.7	0.085	2.35
	2.0	1.01	28.1	0.084	2.37
	5.0 (baseline)	0.97	29.3	0.082	2.41
<i>Panel B. Within-sector demand elasticity</i>					
ε	4.0	0.78	36.3	0.082	2.96
	5.0 (baseline)	0.97	29.3	0.082	2.41
	6.0	1.15	24.6	0.083	2.05

Notes: Panel A varies the elasticity across sector composites, holding the within-sector elasticity and markup at $\varepsilon = 5$ and $\mu = 1.25$. Panel B varies the elasticity across firms within a sector and recomputes the markup $\mu = \varepsilon/(\varepsilon - 1)$ and all cost-share objects in each row, holding $\zeta = 5$. All rows hold $\sigma = 1.6$, the empirical slope $\delta = 1.24$, and the cumulative residualized AI-index return at 22.9%. The GDP weight is γ^{IO} from Proposition 9.

Table G.4: GDP impact under the modified Hulten formula and the exact solution of the model

	Implied productivity news (%)	GDP change, modified Hulten (%)	GDP change, exact (%)
Benchmark, Nov. 2022 – Dec. 2025	32.6	3.61	3.84
Production network, Nov. 2022 – Dec. 2025	29.3	2.41	2.55

Notes: All entries are reported on a log-change basis and multiplied by 100. The first column reports normalized present-value productivity news; the last two columns report corresponding changes in GDP. The first row is the benchmark economy of Section 4.1 and the second row the production network economy of Section 4.3. The implied productivity news is held fixed across the last two columns and is taken from the inversion in the corresponding section; only the map from that news to GDP differs. The modified Hulten column applies the first-order GDP elasticity, γ in the benchmark and γ^{10} in the network economy, to the productivity news. The exact column treats the productivity news as an equivalent log increase, solves the model in levels at the pre-AI and implied productivity levels, and reports the log change in real GDP between the two equilibria. In the production network row, both columns use the same internally consistent production-network calibration described in Appendix G.3; sector output is recomputed so that supply equals demand after capital services and imports are incorporated into the production network. The nonlinear solver assumes that each sector's software engineering payroll share follows a Beta distribution matching its Revelio sales-weighted mean and variance. Calibration is as in the corresponding section: $\varepsilon = 5$, $\sigma = 1.6$, $\delta = 1.24$, and $\mu = 1.25$ in the network economy.

G.5 Proofs

G.5.1 Proof of Lemma 3

Proof. Part (i). Under constant returns and the constant markup, firm (n, j) 's total cost is $\mathcal{R}_n(j)/\mu$, of which the fraction $\alpha_n \gamma_n(j)$ is software engineering compensation. Sector SWE compensation is therefore $\frac{1}{\mu} \int_0^1 \alpha_n \gamma_n(j) \mathcal{R}_n(j) dj$ and sector total cost is $\frac{1}{\mu} \int_0^1 \mathcal{R}_n(j) dj = \mathcal{R}_n/\mu$. Their ratio is

$$\frac{\int_0^1 \alpha_n \gamma_n(j) \mathcal{R}_n(j) dj}{\mathcal{R}_n} = \alpha_n \int_0^1 \varsigma_n(j) \gamma_n(j) dj = \alpha_n \bar{\gamma}_n,$$

using that α_n is common within the sector. The markup cancels, so the sector share is the accounting ratio of SWE compensation to cost for any μ .

Part (ii). Firm (n, j) 's output reaches the economy only as a component of the sector- n composite, of which it supplies the fraction $\varsigma_n(j)$, the same fraction in every transaction, since all buyers purchase the composite. Its total direct-plus-indirect requirement per unit of final expenditure is therefore the sector requirement scaled by its share, $\tilde{\lambda}_n(j) = \tilde{\lambda}_n \varsigma_n(j)$. Hence

$$\int_0^1 \tilde{\lambda}_n(j) \alpha_n \gamma_n(j) dj = \tilde{\lambda}_n \alpha_n \int_0^1 \varsigma_n(j) \gamma_n(j) dj = \tilde{\lambda}_n \alpha_n \bar{\gamma}_n,$$

and summing over n gives part (ii). □

G.5.2 Proof of Lemma 4

Proof. Throughout, $\tilde{w}_t = w_t/S_t$ is the effective SWE wage; we normalize $d \log S_t = 1$, so $u = d \log \tilde{w}_t$ and $v = d \log W_t$.

Part (i): relative price. By Shephard's lemma on the labor composite, $d \log q_n(j) = \gamma_n(j)u + (1 - \gamma_n(j))v$. With the Cobb-Douglas outer nest and the constant markup,

$$d \log p_n(j) = \alpha_n [\gamma_n(j)u + (1 - \gamma_n(j))v] + \sum_m \Omega_{nm} d \log P_m.$$

The composite price moves with the sales-weighted average of individual variety prices, $d \log P_n = \int_0^1 \varsigma_n(j) d \log p_n(j) dj$, which takes the same form with $\bar{\gamma}_n$ in place of $\gamma_n(j)$. Subtracting, the intermediate-price terms cancel, and the two labor terms combine through $\alpha_n[(1 - \gamma_n(j)) - (1 - \bar{\gamma}_n)] = -\alpha_n(\gamma_n(j) - \bar{\gamma}_n)$, giving part (i).

Part (ii): profit deviation. Because every buyer allocates across sector- n varieties with elasticity ε , firm sales are $\mathcal{R}_n(j) = (p_n(j)/P_n)^{1-\varepsilon} \mathcal{R}_n^{\text{tot}}$, where $\mathcal{R}_n^{\text{tot}}$ is total spending on sector n . Hence

$$d \log \mathcal{R}_n(j) = (1 - \varepsilon)(d \log p_n(j) - d \log P_n) + d \log \mathcal{R}_n^{\text{tot}},$$

and the second term is common to all firms in the sector. Profits are the constant fraction $1/\varepsilon$ of sales, so subtracting $d \log E_t$ from the firm sales response, using part (i), and defining $b_n = d \log \mathcal{R}_n^{\text{tot}} - d \log E_t$ gives (G.1). For the stacked form of b , total spending on sector n is final expenditure plus intermediate purchases, $\mathcal{R}_n = f_n E_t + \sum_m \Omega_{mn} \mathcal{R}_m / \mu$, where $E_t = \sum_m P_{mt} C_{mt}$ is nominal final expenditure and intermediate purchases equal cost shares times cost $\mathcal{R}_m / \mu$. Differentiating with constant (Ω, μ) and using $f_n E_t / \mathcal{R}_n + \sum_m \mathcal{T}_{nm} = 1$ gives

$$d \log \mathcal{R}_n - d \log E_t = \frac{f_n E_t}{\mathcal{R}_n} (1 - \zeta) \left(d \log P_n - \sum_m f_m d \log P_m \right) + \sum_m \mathcal{T}_{nm} (d \log \mathcal{R}_m - d \log E_t).$$

Here $\mathcal{T}_{nm} = \Omega_{mn} \mathcal{R}_m / (\mu \mathcal{R}_n)$. Stacking sectors and substituting the price response $d \log P = \Lambda(\alpha \bar{\gamma} u + \bar{\theta} v)$ from part (i) yields $b = \mathcal{B}(\alpha \bar{\gamma} u + \bar{\theta} v)$ with $\mathcal{B} = \Psi \text{diag}(f_n E_t / \mathcal{R}_n) (1 - \zeta) (\mathbf{I} - \mathbf{1} f^\top) \Lambda$ and $\Psi = (\mathbf{I} - \mathcal{T})^{-1}$.

Part (iii): factor-market clearing. SWE labor clears when $w_t s = \frac{1}{\mu} \sum_n \int_0^1 \alpha_n \gamma_n(j) \mathcal{R}_n(j) dj$. Differentiating, two within-sector channels operate: the cost-share response of the labor composite, $d \log(\alpha_n \gamma_n(j)) = (1 - \gamma_n(j))(1 - \sigma)(u - v)$ by Shephard's lemma on the CES nest, and the full sales response of part (ii), which splits into $d \log E_t$, the sector-common b_n , and the deviation $(1 - \varepsilon)\alpha_n(\gamma_n(j) - \bar{\gamma}_n)(u - v)$. Aggregating each channel against SWE compensation uses the two moment identities

$$\int_0^1 \varsigma_n(j) \gamma_n(j) (1 - \gamma_n(j)) dj = \bar{\gamma}_n (1 - \bar{\gamma}_n) - \sigma_{\gamma,n}^2, \quad \int_0^1 \varsigma_n(j) \gamma_n(j) (\gamma_n(j) - \bar{\gamma}_n) dj = \sigma_{\gamma,n}^2.$$

Dividing the SWE market-clearing condition by $w_t s$ and using $d \log s = 0$ yields the first equation below, where v_n are the SWE compensation weights. Applying the same steps to non-SWE market clearing yields the second equation, where v_n^L are the non-SWE compensation weights, the cost-share responses are complementary, and the firm-specific sales deviation is aggregated with weight

$-\sigma_{\gamma,n}^2/(1-\bar{\gamma}_n)$:

$$\begin{aligned} u+1 &= d\log E_t + \sum_n v_n \left[b_n + (1-\bar{\gamma}_n)(1-\sigma)(u-v) \right] + \sum_n v_n [(1-\varepsilon)\alpha_n - (1-\sigma)] \frac{\sigma_{\gamma,n}^2}{\bar{\gamma}_n} (u-v), \\ v &= d\log E_t + \sum_n v_n^L \left[b_n - \bar{\gamma}_n(1-\sigma)(u-v) \right] - \sum_n v_n^L [(1-\varepsilon)\alpha_n - (1-\sigma)] \frac{\sigma_{\gamma,n}^2}{1-\bar{\gamma}_n} (u-v), \\ 0 &= f^\top \Lambda (\alpha \bar{\gamma} u + \bar{\theta} v) = \tilde{s}_S u + \tilde{s}_L v. \end{aligned}$$

The first two equations are the factor-market-clearing conditions, and the last follows from the maintained numeraire $P_t = 1$. The distributions enter only through $(\bar{\gamma}_n, \sigma_{\gamma,n}^2)$. Subtracting the first two equations eliminates $d\log E_t$. All terms in the differenced factor condition depend only on $u - v$. In particular, b is unaffected by a common shift of (u, v) because $\alpha_n \bar{\gamma}_n + \bar{\theta}_n = \alpha_n$ and $\mathcal{B}\alpha = 0$, using

$$(\mathbf{I} - \mathbf{1}f^\top)\Lambda\alpha = (\mathbf{I} - \mathbf{1}f^\top)\mathbf{1} = 0.$$

The differenced factor condition and the numeraire condition therefore form a 2×2 linear system in (u, v) .

One-sector roundabout closed form. With $N = 1$ and $\Omega_{11} = 1 - \alpha$, the centering matrix $(\mathbf{I} - \mathbf{1}f^\top)$ maps the single sector's price response to zero, so $b = 0$. Differencing the first two equations eliminates $d\log E_t$. The coefficient on $u - v$ in the resulting condition is the sum of the cost-share term $(1-\sigma)[1 - \sigma_\gamma^2/(\bar{\gamma}(1-\bar{\gamma}))]$ and the sales-deviation term $\alpha(1-\varepsilon)[\sigma_\gamma^2/(\bar{\gamma}(1-\bar{\gamma}))]$. Collecting terms and using $u - v = (d\log w_t - d\log W_t) - d\log S_t$,

$$u - v = -\frac{1}{\Sigma}, \quad \Sigma = \sigma + [\alpha(\varepsilon - 1) - (\sigma - 1)] \frac{\sigma_\gamma^2}{\bar{\gamma}(1-\bar{\gamma})},$$

which is the closed form of Section 2. □

G.5.3 Proof of Proposition 8

Proof. Step 1: returns. For each aggregate shock, the argument of Lemma 4 applies verbatim: the only firm-level heterogeneity inside a sector is the payroll share, and demand weights $\omega_n(Z_t)$ are sector-level objects, so net of the common expenditure response a shock Z_k generates the profit response $b_n^{Z_k} + (1-\varepsilon)\alpha_n(\gamma_n(j) - \bar{\gamma}_n)(u_k - v_k)$, where (u_k, v_k) are the wage responses to a unit Z_k shock. Capitalizing profits into values and taking period- t revisions exactly as in Lemma 1 yields (G.2).

Step 2: the relative wage does not respond to Z_k . Applying the factor-market-clearing equations in the proof of Lemma 4 to a unit Z_k shock and subtracting the non-SWE condition from the SWE condition gives

$$K_w(u_k - v_k) = \sum_n (v_n - v_n^L) b_n^{Z_k},$$

where

$$K_w = 1 + (\sigma - 1) \sum_n \left\{ v_n(1 - \bar{\gamma}_n) + v_n^L \bar{\gamma}_n - \left(\frac{v_n}{\bar{\gamma}_n} + \frac{v_n^L}{1 - \bar{\gamma}_n} \right) \sigma_{\gamma,n}^2 \right\} + (\varepsilon - 1) \sum_n \alpha_n \left(\frac{v_n}{\bar{\gamma}_n} + \frac{v_n^L}{1 - \bar{\gamma}_n} \right) \sigma_{\gamma,n}^2.$$

After collecting terms, the coefficient on $\sigma - 1$ lies between zero and one. Nonnegativity follows from $0 \leq \sigma_{\gamma,n}^2 \leq \bar{\gamma}_n(1 - \bar{\gamma}_n)$. The upper bound follows from the definitions of v_n and v_n^L and weighted concavity. The terms proportional to $\varepsilon - 1$ are nonnegative. Hence $K_w \geq \min\{\sigma, 1\} > 0$. Assumption 5 sets the right-hand side of the first equation to zero. Therefore $u_k = v_k$ and $\delta_k^Z = (1 - \varepsilon)\bar{\alpha}(u_k - v_k) = 0$ for every k .

Step 3: projection. After the first-stage projection, the firm's AI beta equals a sector-common term plus $\kappa(1 - \varepsilon)(u - v)\alpha_n\gamma_n(j)$. The sector fixed effects absorb the sector-common term. By linearity of the second-stage regression and the definition of $\bar{\alpha}$, its coefficient on the payroll share is

$$\delta = \kappa(1 - \varepsilon)(u - v)\bar{\alpha} = \kappa\delta_{\text{in}}^S.$$

This establishes equation (G.4). The no-network closed form follows from the last part of the proof of Lemma 4. $\square$

G.5.4 Proof of Proposition 9

Proof. Step 1: prices. Because S_t enters costs only through the effective SWE wage, the sector price responses propagate through the Leontief inverse: using the sector cost shares of Lemma 3,

$$d \log P_n = \sum_m \Lambda_{nm} [\alpha_m \bar{\gamma}_m d \log \tilde{w}_t + \bar{\theta}_m d \log W_t].$$

Define $\bar{\Gamma}_n = \sum_m \Lambda_{nm} \alpha_m \bar{\gamma}_m$. Since $\Lambda \alpha = \mathbf{1}$, the coefficient on $d \log W_t$ is $1 - \bar{\Gamma}_n$, so $d \log P_n = \bar{\Gamma}_n d \log \tilde{w}_t + (1 - \bar{\Gamma}_n) d \log W_t$. The numeraire $P_t = 1$ imposes $\sum_n f_n d \log P_n = 0$; with $\sum_n f_n \bar{\Gamma}_n = \tilde{s}_S$ and $d \log \tilde{w}_t = d \log w_t - d \log S_t$, this is

$$\tilde{s}_S d \log w_t + \tilde{s}_L d \log W_t = \tilde{s}_S d \log S_t. \quad (\text{G.7})$$

Step 2: GDP accounting. Both factor supplies are fixed and $P_t = 1$, so the factor-income shares $s_{\tau,t}$ satisfy $d \log s_{S,t} = d \log w_t - d \log Y_t$ and $d \log s_{L,t} = d \log W_t - d \log Y_t$, that is, $d \log Y_t = d \log w_t - d \log s_{S,t} = d \log W_t - d \log s_{L,t}$. Taking the $\tilde{s}$ -weighted average of these two identities, whose weights sum to one, and substituting (G.7),

$$d \log Y_t = (\tilde{s}_S d \log w_t + \tilde{s}_L d \log W_t) - \sum_{\tau} \tilde{s}_{\tau} d \log s_{\tau,t} = \tilde{s}_S d \log S_t - \sum_{\tau} \tilde{s}_{\tau} d \log s_{\tau,t},$$

which is (G.5) with $\tilde{s}_S = \sum_n \tilde{\lambda}_n \alpha_n \bar{\gamma}_n$.

Step 3: when the reallocation term vanishes. At $\mu = 1$, cost equals sales, so the cost-based Do-mar weights solve the same fixed point as sales shares: goods-market clearing gives $\mathcal{R}_n = f_n Y_t + \sum_m \Omega_{mn} \mathcal{R}_m$, hence $\mathcal{R}_n / Y_t = \tilde{\lambda}_n$, and nominal GDP equals factor income, $Y_t = w_t s + W_t L$. Then

$\tilde{s}_\tau = s_{\tau,t}$ and

$$\sum_\tau \tilde{s}_\tau d \log s_{\tau,t} = (s_S d \log w_t + s_L d \log W_t) - d \log Y_t = d \log(w_t s + W_t L) - d \log Y_t = 0,$$

using $P_t = 1$. When $\Omega = 0$, $\alpha_n = 1$, $\Lambda = \mathbf{I}$, and all sector sales are final, so $\mathcal{R}_{nt}/Y_t = f_{nt}$. Hence

$$\begin{aligned} s_{S,t} &= \frac{1}{\mu} \sum_n f_{nt} \tilde{\gamma}_{nt}, \\ s_{L,t} &= \frac{1}{\mu} \sum_n f_{nt} (1 - \tilde{\gamma}_{nt}). \end{aligned}$$

Since $\sum_n f_{nt} = 1$, $s_{S,t} + s_{L,t} = 1/\mu$ at every date. At the steady state, $f_n = \tilde{\lambda}_n$, so these shares equal $\tilde{s}_S/\mu$ and $\tilde{s}_L/\mu$, respectively, and $\tilde{s}_\tau = s_{\tau,t}/(s_{S,t} + s_{L,t})$. Evaluating differentials at the steady state gives

$$\sum_\tau \tilde{s}_\tau d \log s_{\tau,t} = d \log(s_{S,t} + s_{L,t}) = d \log(1/\mu) = 0.$$

Away from these two cases, the reallocation term is computed from the solved system of Lemma 4(iii), with $d \log Y_t$ obtained from the income side, including the profit response.

Step 4: best linear predictor. Multiplying (G.5) by $(1 - \beta)\beta^k$, summing over horizons, and applying the revision operator $(\mathbb{E}_t - \mathbb{E}_{t-1})$ gives $\mathcal{Y}_t = \gamma^{IO} \mathcal{S}_t$. By Proposition 8, $\mathbb{E}^*[\mathcal{S}_t | \tilde{R}_t^{\text{AI}}] = \kappa \tilde{R}_t^{\text{AI}}$, and linearity of the projection gives (G.6). $\square$

H Solving the Model Exactly and Nonlinearly

The modified version of Hulten's theorem, which we use in Section 4.1 and again in Section 4.3, is a first-order result. The gains we infer from asset prices are large enough to warrant checking the accuracy of this approximation. This appendix reports what the model delivers when we do not approximate.

The exercise takes the inferred (log-point) increase in software engineering productivity as given and asks how much GDP it produces. We solve for the equilibrium of the economy twice, once at the pre-AI level of software engineering productivity and once after a gain of the size our inversion implies, and compare real GDP across the two. Both are full solutions of the model in levels: all prices and quantities adjust, profits are rebated to the household, and every market clears. In the network economy, each sector's firms follow a Beta distribution of software engineering payroll shares, with parameters uniquely determined by the Revelio sales-weighted mean and variance. Both the linearized calculation and the level solution use the same internally consistent economy described in Appendix G.3. Table G.4 collects the results.

Solving exactly raises the GDP impact by roughly six percent of itself, and the reason is the same in the benchmark and in the network economy. GDP responds to a software engineering produc-

tivity gain according to the changing cost-based GDP weight.² A productivity gain makes software engineering labor cheaper in effective terms, and firms respond by using more of it. The cost-based GDP weight therefore rises from 0.111 to 0.125 in the benchmark. In the network economy, the corresponding local GDP elasticity rises from 0.082 to 0.092. The modified Hulten formula applies the pre-AI elasticity to the whole gain, whereas the exact calculation allows it to rise as the gain accumulates.

The correction is mostly about substitution between software engineering and other labor, not about the production network. Its size is nearly the same in the benchmark and network economies despite their different production structures. As effective software engineering labor becomes cheaper, its expenditure share rises inside the labor bundle; cross-sector demand reallocation contributes little to the difference between the first-order and level solutions.

References

Baqaei, David Rezza, and Emmanuel Farhi. 2020. “Productivity and misallocation in general equilibrium.” *The Quarterly Journal of Economics*, 135(1): 105–163.

²This weight is not the direct share of firm costs attributable to software engineering. In the benchmark, the direct cost share is $\alpha\gamma$, while the cost-based GDP weight is γ because the $1/\alpha$ in the cost-based Domar weight cancels the α in the direct cost share.