Automation Under Constraints:

Exchange Rates, Interest Rates, and Firm Investment

Meghana Ayyagari[*], *Vojislav Maksimovic*[†], *Rodimiro Rodrigo*[‡] and *Ariel Weinberger*[§]

This version: December 3, 2025

Abstract: The yen depreciation from 2012–2015 reduced robot prices for U.S. firms. Counterintuitively, financially constrained firms - not unconstrained ones - dramatically increased adoption. We rationalize this with a collateral model: robots are pledgeable assets requiring non-pledgeable intangibles, raising the effective cash invoice for constrained firms and amplifying their elasticities. Linking robot and automation machinery imports to Compustat, we find constrained firms are twice as responsive to exchange rate and interest rate shocks, with capital prices dominating (45-60% larger) borrowing costs. Results reveal a collateral channel reallocating technology adoption toward constrained incumbents, with implications for innovation diffusion and competitive dynamics.

JEL Codes: D22, E22, E51, E52, F31, L10, O14, O33.
Keywords: financial constraints, investment, prices of capital goods, automation, robots, technology adoption, monetary policy, credit channel, foreign exchange

[*]School of Business, George Washington University, Email: ayyagari@gwu.edu
[†]Robert H. Smith School of Business, University of Maryland, Email: maxone@umd.edu
[‡]School of Business, George Washington University, Email: rodimiro.rodrigo@gwu.edu
[§]School of Business, George Washington University, Email: aweinberger@gwu.edu

1 Introduction

When the Japanese yen depreciated sharply in 2012–2015, imported industrial robots became substantially cheaper for U.S. firms. Financially constrained firms—whether defined by low cash, debt-constraint language in 10-Ks, or high leverage—scaled up automation markedly more than their unconstrained peers. This amplified response among constrained firms stands in contrast to the standard view in corporate finance that financing frictions dampen constrained firms' ability to scale into favorable shocks.

Most empirical work on financial constraints studies how firms respond to variation in discount rates—policy rates, credit spreads, or bank-lending shocks[1]—or to shocks in internal funds.[2] By contrast, we know far less about how capital-input prices themselves enter the financing margin, despite their importance for new technologies. We study these two components of the user cost[3] jointly, namely capital-goods prices and external financing costs, in a setting where each moves sharply and independently. Two plausibly exogenous shocks provide this variation: (i) yen movements that shift robot prices, and (ii) high-frequency monetary policy surprises that shift financing costs.

Industrial robots have a distinctive financing structure. The hardware—robot arms, controllers, and sensors—is standardized and redeployable, with active secondary markets, making it pledgeable collateral in the sense of Benmelech et al. (2022). Yet converting this hardware into productive capacity requires substantial non-pledgeable intangibles such as software customization, systems integration, worker training, and process redesign, as documented in Rodrigo (2023). This bundling of pledgeable and non-pledgeable components distinguishes robots from both traditional CAPX (mostly pledgeable and low in intangible intensity) and the intangible-heavy investments (predominantly non-pledgeable) emphasized by Falato et al. (2022) and the broader superstar/duration literature (e.g., Kroen et al. (2024), Liu et al. (2023)). Empirically, robots offer an exceptionally well-defined setting. Their concentrated global supply, clean HS classifications, and near-mechanical yen–dollar

[1] See Bernanke and Gertler (1995); Gilchrist and Zakrajšek (2012); Ottonello and Winberry (2020); Gopinath et al. (2017)

[2] See the investment–cash-flow literature and structural financing models, e.g., Fazzari et al. (1988); Whited (1992); Hennessy et al. (2007); Almeida and Campello (2007).

[3] We use "user cost" to refer to the two components of the capital user cost most relevant for imported automation equipment—its purchase price and the external financing cost required to acquire it.

pass-through enable us to trace firm-level user-cost shocks with far greater precision than is possible for aggregate capital expenditure.

To rationalize the empirical patterns, we develop a simple collateral model that captures the financing structure of robot adoption. Firms finance automation with internal cash plus collateralized borrowing, but once they reach their pledge cap they are constrained by their own liquidity. The model collapses all elements of the robot acquisition into a single "invoice": the sticker price plus non-pledgeable costs minus the financed share. Both exchange-rate movements and interest-rate changes operate only through this invoice. Consequently, elasticities are governed by the share of the invoice that must be financed out of cash—rather than by production-function curvature—and the model delivers closed-form mappings from balance-sheet primitives to firms' price and rate sensitivities. The model provides a sharp empirical diagnostic: if liquidity binds, firms with tighter balance sheets must exhibit simultaneously larger price elasticities and larger rate elasticities, a joint pattern that no alternative mechanism generates.[4]

We take the mechanism to the data by assembling a novel firm–year panel that links transaction-level bills of lading for automation imports (2007–2019) to Compustat data and several measures of financial constraints. Two shocks provide plausibly exogenous variation along the two margins of the user cost. First, the rapid *yen depreciation* against the dollar in 2012 - driven by changes in Japanese domestic policy - passed directly through to U.S. robot import prices.[5] Second, we isolate the financing margin by instrumenting the one-year Treasury yield with cumulative high-frequency FOMC surprises following Bauer and Swanson (2023). Together, these shocks allow us to trace how collateral frictions condition investment elasticities to capital-price and interest-rate shocks.

The data line up exactly with this prediction. Our main findings are as follows: First, a 10% yen depreciation raises the unconditional probability that a listed manufacturing firm begins importing robots by 2.25 percentage points (pp). The effect is 1.4–1.8 pps larger among firms classified ex ante as financially constrained—whether by below-median cash, "debt-constraint" language following

[4]This elasticity-based approach is useful more generally: examining how investment or cash-flow sensitivities vary with firm fundamentals can help identify which underlying frictions drive investment and cash-management decisions in other settings.

[5]Japan supplies more than half of U.S. industrial-robot imports, so yen–dollar moves pass straight through to equipment prices; see Table 3, Panel B.

Linn and Weagley (2024) and Hoberg and Maksimovic (2015), or by leverage. Intensive-margin responses in shipment volumes and asset-scaled imports are of similar magnitude.

Second, a 25 basis point easing in the instrumented one-year Treasury rate increases robot adoption 2 pps more among constrained firms, while unconstrained firms respond weakly.

Third, price channels dominate financing channels. In horse-race specifications with both shocks entered simultaneously, exchange-rate effects are 45–60% larger than interest-rate effects across firm types. Input-price variation therefore appears to be a more powerful driver of investment responses than previously recognized in the corporate-finance literature, with implications for how global supply chains and exchange-rate movements shape technological diffusion.

Fourth, the ordering of elasticities is exactly as the model predicts. Financially constrained firms respond strongly to both robot-price shifts and interest-rate movements; unconstrained firms respond little. Across all measures of financing capacity, tighter balance sheets are associated with larger price and rate elasticities. A simple counterfactual underscores the quantitative importance of this mechanism: aggregate robot adoption would have been roughly 17% lower in the absence of the financing channel.

Fifth, event-study and IV local-projection estimates show that the adoption gap opens within one year of a shock and persists for at least four years, ruling out simple intertemporal substitution. Finally, the results survive controls for broad-dollar movements, placebo tests using CAPX and R&D, controls for Tobin's Q and profitability, and conditioning on firm-level outsourcing intensity. These checks confirm that the effects are specific to automation and to the yen- and policy-rate shocks we exploit.

Taken together, these estimates, identified from unexpected components of price and rate movements, highlight a collateral channel through which macro shocks reallocate technology adoption. When automation equipment becomes cheaper or borrowing costs fall, financially constrained incumbents, not unconstrained leaders, emerge as the marginal adopters.

The distributional implications are substantial. Constrained, often lower-productivity, firms gain disproportionate access to productivity-enhancing technologies. This reshapes the pattern of techno-

logical diffusion, alters competitive dynamics, and may amplify inequality across firms and regions. In this sense, shocks to exchange rates, interest rates, and trade policy[6] do more than shift aggregate investment—they reorder who adopts technology. Policymakers seeking to influence not only the aggregate pace but also the distribution of technological adoption should therefore consider this collateral channel alongside traditional demand and discount-rate mechanisms.

A related literature studies robots through their effects on labor demand and productivity. Graetz and Michaels (2018) and Acemoglu and Restrepo (2020) document sizable labor-market and efficiency consequences, while Benmelech and Zator (2025) show that robot investment remains small in aggregate—under 0.3% of equipment spending—and concentrated in a few manufacturing industries, often expanding in response to labor scarcity. Our focus differs: we abstract from labor dynamics and instead study the financial sorting margin, showing that shocks to capital-input prices and financing costs generate systematically different adoption responses across firms with different balance-sheet positions.

Overall, this paper makes three contributions. First, on the empirical side, we document that financially constrained firms are more, not less, responsive to capital-price and financing shocks in automation adoption.

Second, on the theoretical side, we formalize this intuition in a two-period collateral model that delivers closed-form adoption elasticities. By collapsing the adoption problem to an invoice identity, the model links firm primitives directly to investment responses without requiring production-function estimation. Building on the collateral-constraints tradition of Almeida and Campello (2007) and Rampini and Viswanathan (2010), our setup explains why financially marginal adopters of collaterizable automation technology are especially sensitive to shocks in user costs.

Third, we show that diffusion through financially constrained firms scales with the magnitude and frequency of unexpected declines in the user cost of pledgeable capital, helping explain why automation diffusion often begins with financially fragile firms.

Our collateral-amplification mechanism generates risk-sorted investment elasticities that mirror,

[6] See Reuters, "US opens tariff probes into medical equipment, robotics, industrial machinery," September 24, 2025.

but reverse, the superstar/duration channel of Kroen et al. (2024): whereas their evidence shows that in a low-rate environment falling interest rates disproportionately benefit industry leaders with long-duration intangible assets, our model identifies the opposite regime. Constrained followers are the marginal adopters of collaterizable automation when financing frictions bind. The two channels operate in different asset regimes - intangibles versus pledgeable capital - and can coexist.

The remainder of the paper is organized as follows: Section 2 situates our contribution within the literature on financing frictions, exchange-rate pass-through, and automation; Section 3 lays out the collateral model and derives closed-form elasticities; Section 4 describes the novel bill-of-lading dataset and constraint measures; Section 5 documents baseline adoption patterns and estimates the causal effects of yen and rate shocks. Section 6 maps the results back to the model, and Section 7 concludes.

2 Literature

Our work connects three strands of research: financing frictions and investment, collateral and pledgeability, and heterogeneous transmission of price and macro shocks to firm investment. The classic view is that financial frictions prevent constrained firms from scaling into profitable opportunities.[7] With borrowing frictions or scarce collateral, constrained firms expand investment less when borrowing costs fall or favorable policy/demand shocks arrive, relative to unconstrained peers. Foundational evidence comes from investment-cash flow sensitivity tests (Fazzari et al., 1988) and structural models (Whited, 1992; Hennessy et al., 2007; Farre-Mensa and Ljungqvist, 2016). Shock-based studies, likewise, find muted responses among constrained firms to bonus depreciation (Edgerton, 2010) and to monetary easing (Ottonello and Winberry, 2020). In contrast, we show that constrained firms can experience growth spurts when the user cost of pledgeable capital falls (via prices or rates). Both shocks operate by relaxing cash constraints for the marginal adopters of automation technology. More constrained firms, therefore, exhibit the largest (in magnitude) elasticities to both price and rate shocks.[8]

[7] For a review, see Frank and Goyal (2009).

[8] This interpretation complements early accelerator frameworks with collateral constraints by Kiyotaki and Moore (1997), Bernanke et al. (1999), and Holmstrom and Tirole (1997).

Our model builds on the financial constraints literature that emphasizes the dual role of cash holdings and asset pledgeability in corporate investment. It is closest in spirit to Rampini and Viswanathan (2010, 2013), who emphasize that pledgeability is an endogenous choice and that constrained firms conserve debt capacity rather than fully collateralizing assets, with collateralization itself consuming scarce resources. We extend this logic to automation adoption by collapsing all costs into a unit-level "invoice" that combines pledgeable purchase price with non-pledgeable intangible capital. Earlier work by Almeida et al. (2004); Almeida and Campello (2007) highlights the importance of liquidity management and the collateral multiplier, while Bolton et al. (2011) bring these insights into dynamic settings. Our two-period framework preserves the essential trade-offs identified in this literature while remaining parsimonious enough to generate sharp, testable predictions about adoption elasticities.

Recent work has refined our understanding of how collateral constraints bind: Lian and Ma (2021) document that 80% of US corporate lending is cash flow-based rather than asset-based, although asset-based constraints become more binding during downturns; Drechsel (2023) further show that these earnings-based borrowing constraints amplify monetary policy transmission, with firms facing tighter earnings-based constraints exhibiting stronger investment responses to interest rate changes—a mechanism consistent with our model's prediction that the financing share determines rate sensitivity; Chaney et al. (2012) and Catherine et al. (2022) show that real estate collateral values significantly affect firm investment through a borrowing channel; and Benmelech et al. (2022) demonstrate that asset redeployability—not just tangibility—determines debt capacity. The rising importance of intangible capital documented by Falato et al. (2022) makes our distinction between pledgeable equipment and non-pledgeable intangible capital increasingly relevant. Our contribution is to show that even within a single asset class—robots—collateralization depends not just on asset values but on the per-unit split between purchase price and non-pledgeable intangibles, creating a new margin of adjustment that drives heterogeneous elasticities to price and rate shocks.

Robots thus provide a sharp laboratory for studying pledgeable-intangible interactions. The hardware—robot arms, controllers, sensors—is standardizable, redeployable capital with active secondary markets, fitting the secured-lending perspective of Benmelech et al. (2022). Yet profitability turns critically on complementary intangibles: software, integration, training, and process redesign, con-

sistent with Falato et al. (2022). Our model formalizes this hybrid structure by embedding both a pledgeable purchase price and a non-pledgeable per-unit intangible cost in the same adoption "invoice," and our empirics exploit precisely this mix.

Our model also relates to work showing that falling rates can favor large, intangible-intensive "superstars" via duration and easy finance (Kroen et al., 2024; Liu et al., 2023). In our pledgeable-capital regime, the cross-section reverses: reductions in the policy rate stimulate capital spending more among collateral-dependent, high-spread firms than among low-spread firms. This mechanism applies to the vast majority of non-leader firms and is orthogonal to the superstar/duration channel emphasized in intangible-heavy settings. It complements evidence on heterogeneity in monetary transmission by age, liquidity, and financial structure (Cloyne et al., 2023; Ottonello and Winberry, 2020; Jeenas, 2024).

Our contribution extends this literature by examining a specific type of capital investment—automation technology—typically undertaken by large, mature firms traditionally thought to be unconstrained. However, conditional on size and age, we show that leverage and debt constraints still play important roles. Our theoretical insight is that riskier firms exhibit greater sensitivity to capital price shocks through an endogenous financial accelerator mechanism.

Existing research has primarily examined automation's consequences once adopted, including labor market effects (Acemoglu and Restrepo (2020), Benmelech and Zator (2025), Dixon et al. (2021)) and productivity (Graetz and Michaels (2018); Koch et al. (2021); Rodrigo (2023)). However, significantly less attention has been paid to the critical first-order question: which firms adopt automation technologies and why? Our research addresses this gap by showing how macroeconomic conditions, specifically exchange rates and interest rates, create uneven patterns of adoption that may amplify existing firm-level advantages rather than democratizing access to productivity-enhancing technologies.[9]

Our findings align with micro evidence on firms' stated barriers. In the American Business Survey 2019 analyzed by Acemoglu et al. (2022), firms most frequently cite that robots are "too expensive"

[9]See also Wang (2022), who measures robot adoption using LFTTD imports.

and that they face a "lack of capital" as reasons for not adopting. These map directly to our model's key frictions: a high invoice (price plus intangibles) and tight cash constraints, reinforcing that adoption is jointly gated by prices and internal liquidity rather than technological suitability alone.

Measurement and pass-through. We study a specific class of tangible investment—automation technology—that is both partially pledgeable and heavily imported. We measure adoption from transaction-level bills of lading (BoL) linked to Compustat/10-K text, which delivers firm-level timing on the *adoption margin* (France/Denmark comparables include Bonfiglioli et al. 2024; Acemoglu et al. 2020; Humlum 2022; see also U.S. survey/IFR evidence Bessen et al. 2025; Koch et al. 2021; Antonioli et al. 2024; Brynjolfsson et al. 2023). Our identification leverages two sources: (i) *exchange-rate pass-through* to imported input prices—robots are largely priced in yen—so yen movements shift the dollar purchase price directly (e.g., Amiti et al., 2014; Blaum et al., 2018; Gopinath and Neiman, 2014; Di Giovanni et al., 2022); and (ii) *high-frequency monetary policy surprises*, following Döttling and Ratnovski (2023) in using the one-year Treasury yield and employing the orthogonalized HFI series of Bauer and Swanson (2023). We use the orthogonalized series precisely to mitigate Fed-information confounds emphasized by Gertler and Karadi (2015); Jarociński and Karadi (2020); Miranda-Agrippino and Ricco (2021). Whereas prior work emphasizes intangibles/high-q firms, we show that automation investment—which is intensive in start-up costs—is also rate-sensitive, but through a liquidity channel. We document that the effects are *automation-specific* through CAPX and R&D placebo exercises.

3 Model

We develop a two-period model that links firms' robot-investment elasticities directly to balance-sheet primitives by characterizing how much internal cash each unit of capital requires. The model is designed to: (i) abstract from production or demand estimation by focusing on financing feasibility when constraints bind; (ii) express the response to price and rate shocks through a transparent sources–uses identity; and (iii) align naturally with the short-run exogenous shocks we exploit (exchange-rate pass-through to p_k and high-frequency monetary surprises to R).

3.1 Economic Environment

Time is $t \in \{0, 1\}$. At $t = 0$ a firm with internal cash c chooses robot capital k', borrowing b', and an encumbrance share $\theta \in [0, \bar{\theta}]$. Robots have unit purchase price p_k and require complementary intangible capital h - the software, integration, training, and process-adaptation expenditures needed to make the hardware productive. Industry evidence indicates that integration and related expenditures are at least as large as the hardware cost, and often substantially larger.[10] We therefore treat h as a per-unit non-pledgeable component, capturing the dual structure of robot investment: pledgeable hardware bundled with substantial non-pledgeable intangibles. Scaling entails convex installation cost $\frac{\tau}{2}(k')^2$.

At $t = 1$, uncertainty is realized. With survival probability π, output equals $A(k')^\alpha$ and the undepreciated resale value is $(1 - \delta)p_k k'$. With probability $1 - \pi$, the firm defaults and the asset liquidates for $\rho p_k k'$, with $\rho \in (0, 1)$. The loan contract gives creditors a legal claim on a fraction θ of the asset, but because of enforcement frictions, they recover only $m(\theta) \leq \theta$. Equity forfeits the pledged share θ, receiving $(1 - \theta)\, \rho p_k k'$ in default; creditors receive $m(\theta)\, \rho p_k k'$; and the residual dissipates as a deadweight loss. The encumbrance choice θ therefore determines *effective pledgeability* via an increasing, concave mapping: $m : [0, \bar{\theta}] \to [0, \bar{m}]$ with $m(0) = 0$, $m'(\theta) > 0$, and $m''(\theta) < 0$, where $\bar{m} \equiv m(\bar{\theta})$.

3.2 Debt contract, pricing, and borrowing capacity

Lenders are risk–neutral and discount at the gross policy rate $R = 1 + r_b$. The loan is one-period and non-recourse. At $t = 1$, with probability π ("survival"), the borrower repays b'; with probability $1 - \pi$ ("default"), the lender seizes the encumbered share θ of the asset, but due to enforcement frictions, recovers only $m(\theta)\, \rho\, p_k\, k'$.

[10] Leigh and Kraft (2018) document that local robot integrators account for roughly three-quarters of total installed system costs, a fact noted in the online appendix of Acemoglu and Restrepo (2020). Engineering trade sources likewise emphasize that installation, programming, safety configuration, and line re-engineering materially increase total expenditures beyond the hardware price; see Engineering.com (2018), "Time and Money: How Much Do Industrial Robots Cost?"

Under competition, the date–0 price equals the expected discounted payoff:

$$\frac{\pi\, b' + (1-\pi)\, m(\theta)\, \rho\, p_k\, k'}{R} \;=\; b'. \tag{3.1}$$

Solving for the ratio of borrowing to effective collateral, $b'/(m(\theta)p_k k')$, yields the advance rate per dollar of effective liquidation–value collateral:

$$y(R,\pi,\rho) \;=\; \frac{(1-\pi)\,\rho}{R-\pi}, \qquad R>\pi, \tag{3.2}$$

so that $b' = y(R,\pi,\rho)\, m(\theta)\, p_k\, k'$ and the borrowing constraint can be written as in (3.3):[11]

$$b' \;\le\; y(R,\pi,\rho)\, m(\theta)\, p_k\, k',\,[12] \tag{3.3}$$

subject to the hard cap on encumbrance

$$0 \;\le\; \theta \;\le\; \bar{\theta} \qquad \text{(equivalently, } 0 \le m(\theta) \le \bar{m}\text{)}. \tag{3.4}$$

The linear special case is $m(\theta)=\theta$, which reduces (3.3) to $b' \le y\,\theta\, p_k\, k'$.[13]

Regularity and monotone effects. We maintain $0<\pi<1$, $0<\rho<1$, $R>1$, and $R>\pi$. Under these assumptions, y in (3.2) is well–defined, strictly decreasing in R, and strictly increasing in ρ. These monotone properties are used later when we study how interest–rate and price shocks shift the feasible set.

[11] We will use $\chi_R \equiv \partial \ln y/\partial \ln R = -\frac{R}{R-\pi}$ below.

[12] Large US firms often borrow under both asset-based and cash-flow-based arrangements in the sense of Lian and Ma (2021). Below we show that our borrowing constraint can be interpreted as a reduced-form financing frontier that nests both types of contracts; see Appendix IA.3.14 for a detailed discussion.

[13] We measure b' in date–0 proceeds (cash raised). If a loan with face value B' trades at price $q(R,\pi,\rho)$, then $b' \equiv qB'$. Under (3.1)–(3.2), the collateral haircut is captured by $m(\theta)$, while $y(R,\pi,\rho)$ is the price per dollar of effectively recoverable collateral.

3.3 Budget identity and feasibility

Robot acquisition at $t = 0$ requires upfront cash outlays financed by internal funds and borrowing. The *sources–uses* identity is

$$\underbrace{c + b'}_{\text{sources at } t=0} \;=\; \underbrace{(p_k + h)\, k' + \tfrac{\tau}{2}\,(k')^2}_{\text{uses at } t=0}. \tag{3.5}$$

A feasible choice is a triple (k', b', θ) that satisfies the budget identity (3.5), the borrowing frontier (3.3), and the cap (3.4), with $k' \geq 0$, $b' \geq 0$, and $0 \leq \theta \leq \bar{\theta}$.[14]

3.4 Objective Function

Given the environment and constraints, the firm chooses next–period capital k', borrowing b', and pledge share θ to maximize the date–0 present value of equity.

At $t = 1$, if the firm survives (with probability π),[15] it produces $A(k')^\alpha$, keeps the undepreciated share $(1 - \delta)p_k k'$, and repays b' to the lender. Equity's payoff in this state is therefore

$$\Pi_1^{\text{good}} \;=\; A(k')^\alpha + (1 - \delta)p_k k' - b'.$$

If it defaults (with probability $1 - \pi$), the lender seizes the encumbered share $\theta\rho p_k k'$ and equity retains the unencumbered part of the liquidation value, so its payoff in the bad state is

$$\Pi_1^{\text{bad}} \;=\; (1 - \theta)\,\rho\, p_k\, k'.$$

Under competitive, risk–neutral pricing, the loan is priced so that the lender's expected discounted

[14] Throughout we assume parameters are such that each unit of robot capital requires a nonnegative own–cash contribution at $t = 0$. Equivalently, the financed share of the per–unit outlay does not exceed the per–unit purchase plus setup cost: $y(R, \pi, \rho)\, m(\theta) \;<\; 1 + \frac{h}{p_k}$ for all $\theta \in [0, \bar{\theta}]$. This ensures the feasibility set is nonempty and rules out knife–edge cases where the entire outlay could be debt–financed. Detailed consequences of (3.5) and (3.3)—including how they map into the "invoice" representation—are developed in Section 3.4, after we introduce the additional notation needed there.

[15] We take π to be exogenous. At the cost of more algebra, we can generalize the model so that π is a function of θ, with $\pi'(\theta) < 0$. This would leave all the principal results unchanged and enable us to dispense with assumption A(7) in the Appendix.

payoff equals the amount lent:

$$\mathbb{E}_0\left[\frac{\text{repayment}}{R}\right] \;=\; \frac{\pi\, b' + (1-\pi)\, m(\theta)\, \rho\, p_k\, k'}{R} \;=\; b'.$$

From the firm's perspective, the present value of the financing leg is therefore

$$\text{PV(financing)} \;=\; \underbrace{b'}_{\text{inflow at } t=0} \;-\; \underbrace{\mathbb{E}_0\left[\tfrac{\text{repayment}}{R}\right]}_{\text{outflow in PV}} \;=\; b' - b' \;=\; 0.$$

Thus, the loan has zero net present value for the firm: b' affects the objective only through the constraints, not directly via the payoff.

Using this pricing condition, we can drop the b' terms from the payoff and write the firm's problem in terms of the *real* cash flows of the project (output plus resale net of purchase, setup, and installation costs), treating c as given. The firm's problem is

$$\begin{aligned}\max_{k',b',\theta} \quad & \frac{\pi}{R}\Big[A(k')^{\alpha} + (1-\delta)p_k k'\Big] + \frac{1-\pi}{R}\Big[\rho(1-\theta)p_k k'\Big] \\ & - (p_k + h)k' - \tfrac{\tau}{2}(k')^2,\end{aligned} \tag{3.6}$$

subject to the budget identity (3.5), the borrowing–capacity constraint (3.3), and the pledge–cap constraint (3.4).

Fair pricing of debt. As in (3.2), competitive, risk–neutral lenders set the advance rate

$$y(R, \pi, \rho) = \frac{(1-\pi)\rho}{R - \pi}$$

so that

$$\mathbb{E}_0\left[\tfrac{\text{repayment}}{R}\right] = b'.$$

Under this condition the present value of the financing leg is zero, and b' enters the firm's problem only through the budget and capacity constraints.

Let $\lambda \geq 0$ be the multiplier on the budget identity (3.5), $\mu \geq 0$ the multiplier on the borrow-

ing–capacity constraint (3.3), and $\nu^-, \nu^+ \geq 0$ the multipliers on the pledge–cap constraint (3.4). Because the financing leg has zero net present value and b' appears only in the two constraints, the first–order condition with respect to b' implies $\lambda = \mu$, so there is a single shadow value of financial slack (an explicit derivation is given in Appendix IA.3.3). We therefore denote this common value simply by λ and interpret it as the marginal value of one more dollar of financing capacity.

This delivers a distinction between cash–rich and cash–poor firms. When finance is slack, the *cash–rich* case, both the budget and capacity constraints are slack, and both multipliers vanish $\lambda = \mu = 0$; in this case, k' is chosen by the technology FOC (MPK = user cost net of expected resale) and the level of cash does not affect the scale. When finance binds, the *cash–poor* case, $\lambda = \mu > 0$ and the marginal value of one more unit of internal funds equals the marginal value of one more unit of borrowing capacity, and investment is pinned down by the interaction of the budget and collateral constraints: the firm must finance the entire per–unit cash outlay on robots out of its given internal cash c plus what it can borrow against the collateral value of those robots.

3.5 Derivation of the Cash Identity for Cash-poor Firms

Starting from the "Sources = Uses" equality (3.5), which is a binding constraint for these firms, substitute the borrowing–pledge constraint

$$\underbrace{b'}_{\text{borrowing}} = \underbrace{y\, m(\theta)\, p_k\, k'}_{\text{PV on pledge}},$$

and rearrange to obtain the fundamental cash identity:

$$c = (p_k + h - y\, m(\theta)\, p_k)k' + \tfrac{\tau}{2}(k')^2 \tag{3.7}$$

$$= k' \underbrace{\left[p_k(1 + a - y\, m(\theta)) + \tfrac{\tau}{2}k'\right]}_{M}, \tag{3.8}$$

where $a \equiv \frac{h}{p_k}$.

Therefore: $k'M = c$. Define the *internal invoice per unit* M as the own-cash requirement accompanying

each robot purchase:

$$M \equiv p_k(1 + a - y\, m(\theta)) + \tfrac{\tau}{2}k'.$$

We refer to this as the **cash identity**. It compresses all financing frictions into a single reduced form for cash-poor firms: each unit of capital requires an invoice M of internal cash, so scale is pinned by $k' = c/M$ whenever financing binds.

For a cash-poor firm, the act of financing itself is the binding margin. Because the budget and borrowing constraints both bite, every additional unit of capital must be backed one-for-one by the firm's own-cash invoice M. Substituting the binding borrowing condition into the sources–uses identity collapses the full problem into the simple relationship $k'M = c$. The key envelope logic is that we do not need to solve explicitly for how the optimal pledge θ^* adjusts with fundamentals: its first-order condition ensures that this adjustment is captured inside M. As a result, all comparative statics of interest can be obtained by differentiating M directly.

This is why the internal invoice provides a direct bridge from financing frictions to empirical elasticities. When the firm is cash-poor, shocks to the price of capital p_k, the interest rate R, or the non-pledgeable cost h shift the invoice M, and—because scale is pinned by $k' = c/M$—these shifts translate immediately into observable investment responses. In contrast, for a cash-rich firm the budget is slack, $\theta^* = 0$, and M is just accounting rather than a determinant of scale; investment is then chosen by the production technology. Thus, the fact that elasticities can be read off from M is not an odd coincidence, but a direct consequence of constraints binding in equilibrium for cash-poor firms.

Figure 1 decomposes the internal invoice M into the sticker price net of collateralized borrowing $(p_k(1 - y\, m(\theta)))$, the non-pledgeable setup cost h, and the installation term $\frac{\tau}{2}k'$, which we will refer to below as the "exposure" (orange) and "brake" (purple) components.

There are two subcases of cash-poor firms that are similar but give rise to slightly different algebraic analysis below. Cash-poor firms that are "at the cap" are those whose borrowing θ is at the maximum feasible. Cash-poor firms "in the interior" are those for which θ is below the maximum feasible, and can increase leverage in response to shocks.

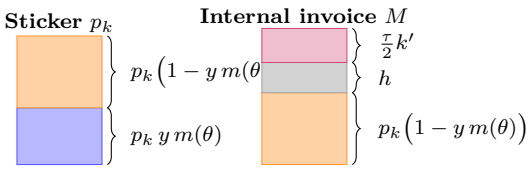


Figure 1: Decomposition of sticker price and definition of the internal invoice M. The first term is the sticker price net of collateralized borrowing, the second term the intangible setup cost, and the third the installation outlay.

Figure 2 shows how scale k' is determined in the three regimes. Panel (a) (financing slack) plots marginal product against user cost; k' is chosen where MPK equals user cost. Panels (b) and (c) (financing binds) plot the cash requirement $k'\,M(\cdot)$ (purple) against available cash c (red); k' is chosen at their intersection. In (b) the pledge choice is interior, so M depends on the optimally chosen θ^*; in (c) the pledge cap binds and M is evaluated at $\bar{\theta}$. The purple curve plots total cash required $k'M(k',\cdot)$; it is increasing and strictly convex for $\tau > 0$.

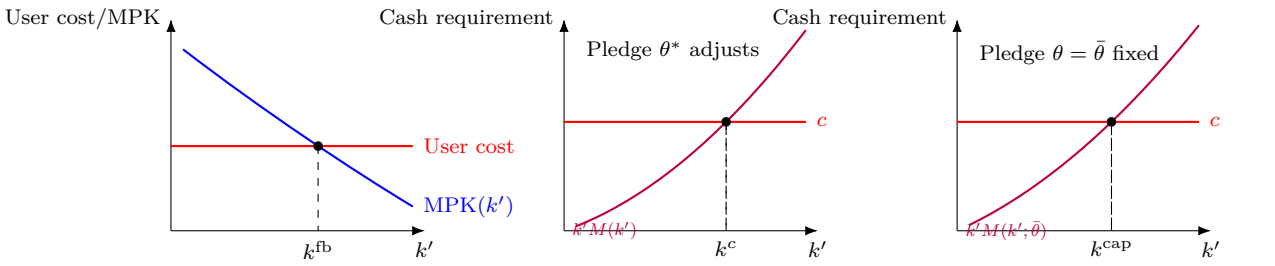


Figure 2: Three regimes. **(a)** Cash-rich (financing slack): choose k' where MPK(k') equals user cost. **(b)** Cash-poor, interior pledge: financing binds; k' solves $k'\,M(k';\theta^*) = c$. **(c)** Cash-poor at the cap: financing binds and $\theta = \bar{\theta}$, so k' solves $k'\,M(k';\bar{\theta}) = c$.

We illustrate the cash-identity logic by deriving the elasticities of k' with respect to p_k and R for cash–poor firms. Take the case where the firm's borrowing capacity is exhausted in equilibrium ($\theta^* = \bar{\theta}$). Borrowing is capped and the budget constraint binds tightly. We call this case the "Cap Regime". Starting with $k'M = c$ in equilibrium, we can derive the price elasticity of capital in three steps. First, log–differentiate the identity: holding c fixed, $d\ln k' + d\ln M = 0$. Second, decompose

$d \ln M$: with $M = p_k\big(1 + a - y\, m(\bar{\theta})\big) + \frac{\tau}{2} k'$ and $a \equiv h/p_k$,

$$d \ln M = \underbrace{\frac{\partial \ln M}{\partial \ln p_k}}_{W_{p_k}} d \ln p_k + \underbrace{\frac{\partial \ln M}{\partial \ln k'}}_{\kappa} d \ln k',$$

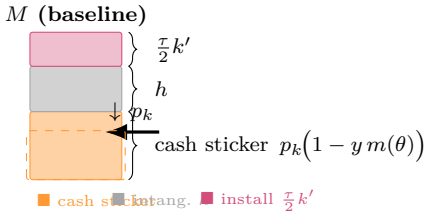
where $W_{p_k} = \frac{p_k\big(1 - y\, m(\bar{\theta})\big)}{M}$ and $\kappa = \frac{\frac{\tau}{2} k'}{M}$. Finally,

$$(1 + \kappa)\, d \ln k' = -\, W_{p_k}\, d \ln p_k \quad \Longrightarrow \quad \boxed{\varepsilon_{k', p_k} = -\frac{W_{p_k}}{1 + \kappa} < 0}.$$

Thus, the elasticity equals the negative ratio of exposure (sticker–net–of–finance share W_{p_k}) to brake (one plus the installation share κ). By the same steps (using $\partial \ln M / \partial \ln R = -W_f\, \chi_R$), we obtain

$$\boxed{\varepsilon_{k', R} = \frac{W_f\, \chi_R}{1 + \kappa} \; < \; 0}, \qquad W_f = \frac{p_k\, y\, m(\bar{\theta})}{M}, \quad \chi_R = \frac{\partial \ln y}{\partial \ln R} < 0.$$

Our diagrammatic decomposition of M can provide some intuition on these derivations. As Figure 1 highlights, a price shock moves only the orange block $p_k(1 - y\, m(\theta))$, while the rate channel operates through y inside that same block.



For a given pledge $m(\theta)$ and scale k', a price shock changes only the orange block:

$$M \;=\; \underbrace{p_k\big(1 - y\, m(\theta)\big)}_{\text{price/finance}} \;+\; \underbrace{h}_{\text{capability}} \;+\; \underbrace{\tfrac{\tau}{2} k'}_{\text{execution}} .$$

This gives a rule of thumb in the cash–poor region:

- A $\downarrow p_k$ shock reduces the net cash sticker only (orange component).
- Greater financing reach (higher y or higher $m(\theta)$) also shrinks the orange block.

- Scale is mechanical: $k' = \dfrac{c}{M}$. Lower $M \Rightarrow$ more units.

Let $w_{\text{price}} \equiv \dfrac{p_k(1 - y\, m(\theta))}{M}$ be the orange share of M. Then, for a given pledge (e.g., at the cap or locally holding $m(\theta)$ fixed),

$$\boxed{\frac{\partial \ln k'}{\partial \ln p_k} \;=\; -\,\frac{w_{\text{price}}}{1+\kappa}, \qquad \kappa \equiv \frac{\frac{\tau}{2}k'}{M}}$$

which shows the response magnitude is proportional to the orange share and attenuated by the installation "brake" $1 + \kappa$.

When the pledge choice is interior ($0 < \theta^* < \bar{\theta}$), the firm optimally adjusts θ in response to shocks, creating an additional margin of adjustment. With a mild regularity condition ruling out the implausible scenario where a firm can fully insulate itself from cash shocks through leverage adjustments the steps are analogous. Proposition 1 summarizes the results. Details appear in the appendix.

Proposition 1 (Sign Preservation of Elasticities in the Interior Regime)**.** *Suppose the firm is cash-poor with an interior pledge choice $0 < \theta^* < \bar{\theta}$. Then the price and rate elasticities of investment satisfy:*

$$\varepsilon^{int}_{k',p_k} \;<\; 0, \qquad \varepsilon^{int}_{k',R} \;<\; 0,$$

where the closed-form solutions are given in the Appendix. Thus, in the interior regime, pledge adjustment can dampen the sensitivity to prices and amplify or mitigate the sensitivity to rates, but in both cases the elasticities remain negative.

Economically, this reflects the fact that only the hardware is pledgeable, while the intangible and installation components must still be paid out of internal funds. Having obtained closed–form solutions for the elasticities of capital with respect to p_k and R in the binding case, we can also determine how their magnitudes depend on the levels of other primitives.

Proposition 2 (Second-Order Cross-Effects with a Level Shock in z)**.** *Assume shocks in $x \in \{p_k, R\}$ are small enough not to switch regimes (interior $0 < \theta^* < \bar{\theta}$ or cap $\theta^* = \bar{\theta}$). In the finance-binding*

region, the first elasticity $E_x \equiv \partial \ln k'/\partial \ln x$ *satisfies*

$$\frac{\partial E_x}{\partial z} \;=\; -\frac{\partial}{\partial z}\Big(\frac{\partial \ln M}{\partial \ln x}\Big) \;=\; -\frac{\partial^2 \ln M}{\partial \ln x\, \partial z}, \qquad \textit{(total derivative along } (k'(x,z),\, \theta^*(x,z))\textit{)}.$$

See Appendix §IA.3.11 for the proof.

In addition to affecting elasticities, shocks to the primitives also affect the endogenous financial variables, leverage θ, and the marginal cost of cash λ. The effects on λ and θ differ, depending on the primitive driving the changes and on whether the firm is at the cap or not.

For cash-poor firms, the b'–FOC gives $\lambda = \mu > 0$: the marginal value of one more unit of internal funds equals the marginal value of one more unit of borrowing capacity. With an interior pledge, θ adjusts when a primitive changes, and the θ–FOC pins down

$$\lambda\, y\, m'(\theta^*) = \frac{1-\pi}{R}\rho \quad \Longrightarrow \quad \lambda = \frac{(1-\pi)\rho}{R\, y\, m'(\theta^*)}.$$

By contrast, at the cap the Kuhn-Tucker (KKT) condition implies the lower bound

$$\lambda \geq \frac{(1-\pi)\rho}{R\, y\, m'(\bar{\theta})}.$$

Hence, at the cap (where θ is fixed) λ rises with cash scarcity, either caused directly by a reduction in c or indirectly by a change in another primitive. Inside the cap, θ adjusts to close the funding gap; the interior θ–FOC determines θ^*. θ^* increases directly with tighter cash or indirectly by a change in another primitive.

These relationships between changes in primitives, changes in elasticities derived above, and changes in θ and λ are useful because, as we show below, observing the predicted relations between primitives, elasticities, θ and λ allows us to test which primitives are driving firm behavior.

In Tables 1 and 2 we summarize how the key primitives $(c, h, \tau, \bar{\theta}, y)$ affect the shadow value of funds λ, the optimal pledge θ, and the elasticities of investment with respect to capital prices and rates. The tables show that the invoice lens cleanly separates the sources of tightness: at the cap,

Table 1: Comparative statics at the cap ($\theta = \bar{\theta}$)

Primitive change	λ	$\lvert\varepsilon_{k',p_k}\rvert$	$\lvert\varepsilon_{k',R}\rvert$	θ
Cash $c \downarrow$ (tighter liquidity)	$\uparrow$	$\uparrow$	$\uparrow$	fixed at $\bar{\theta}$
Intangibles $h \downarrow$ (setup cheaper)	$\downarrow$	$\uparrow$	$\uparrow$	fixed at $\bar{\theta}$
Install frictions $\tau \downarrow$ (easier)	$\downarrow$	$\uparrow$	$\uparrow$	fixed at $\bar{\theta}$
Advance cap $\bar{\theta} \downarrow$ (tighter)	$\uparrow$	$\uparrow$	$\downarrow$	fixed at $\bar{\theta}$
Safer debt $y \uparrow$ ($\pi \uparrow$ or $\rho \uparrow$)	$\downarrow$	$\downarrow$	$\uparrow$	fixed at $\bar{\theta}$

Notes. With pledge fixed at the cap, elasticities follow "exposure ÷ brake":

$$\varepsilon_{k',p_k} = -\frac{W_{p_k}}{1+\kappa}, \quad \varepsilon_{k',R} = \frac{W_f\,\chi_R}{1+\kappa}.$$

Table 2: Comparative statics inside the cap ($0 < \theta^* < \bar{\theta}$)

Primitive change	λ	θ^*	$\lvert\varepsilon^{\text{int}}_{k',p_k}\rvert$	$\lvert\varepsilon^{\text{int}}_{k',R}\rvert$
Cash $c \downarrow$ (tighter liquidity)	0 (diagnostic: θ^*)	$\uparrow$	$\uparrow$	$\uparrow$
Intangibles $h \downarrow$ (setup cheaper)	0 (diagnostic: θ^*)	$\downarrow$	$\uparrow$	$\uparrow$
Install frictions $\tau \downarrow$ (easier)	0 (diagnostic: θ^*)	$\approx$ flat	$\uparrow$	$\uparrow$
Safer debt $y \uparrow$ ($\pi \uparrow$ or $\rho \uparrow$)	0 (diagnostic: θ^*)	$\downarrow$	$\downarrow$	$\uparrow$
Advance cap $\bar{\theta}$ (non-binding locally)	0 (locally)	0 (locally)	0 (locally)	0 (locally)

Notes. Inside the cap, pledge reacts. Elasticities are

$$\varepsilon^{\text{int}}_{k',p_k} = -\frac{W_{p_k} - W_\theta\,\eta_{\theta,p}}{1+\kappa}, \quad \varepsilon^{\text{int}}_{k',R} = \frac{W_f\,\chi_R + W_\theta\,\eta_{\theta,R}}{1+\kappa}.$$

λ is the diagnostic margin, while inside the cap it is θ^*. In both regimes the elasticities retain the same "wedge ÷ brake" structure, making it straightforward to interpret how primitives map into observable sensitivities and to distinguish liquidity, collateral, and cost-side drivers.

3.6 Taking the Model to the Data: Intangible Costs h, Cash c, and the Invoice Lens

We illustrate the use of the invoice lens mechanism by asking the following question. Is it the case that the pattern of elasticities with respect to shocks we observe is consistent with differences in financial constraints, or could they be driven by differences in the efficiency in firms' use of intangible capital?

Robots differ not just in sticker price p_k, but also in complementary intangibles required for integration (e.g., training, customization) h. These costs are paid in cash and cannot be pledged. A higher level of h soaks up internal funds much like a reduction in cash c: both tighten the budget and, for cash-poor firms, tend to increase the magnitudes of the price and rate elasticities (through a larger exposure

share). If h is not directly observed, elasticities alone can therefore confound high h with low c.

The joint movement of λ (at the cap) and θ^* (inside the cap) disentangles these drivers. At the cap, a fall in c raises the shadow value of funds λ, whereas a fall in h lowers or leaves λ unchanged—so $(\lambda, |\varepsilon|)$ moves in opposite directions across the two cases, separating low cash from cheaper intangibles as driver of elasticity changes. In the interior regime ($0 < \theta^* < \bar{\theta}$), the pledge θ adjusts to close the funding gap. Reductions in both c and h tend to raise elasticities, but in opposite ways through θ^*: tighter cash c increases θ^* (the firm encumbers more), whereas cheaper intangibles h reduce θ^* (less need to pledge). Hence, elasticities alone could confound h with c, but the joint movement of θ^* (inside the cap) or λ (at the cap) diagnoses which margin is driving the change.

Thus, the pattern of three observable objects – the two elasticity magnitudes $|\varepsilon_{k',p_k}|$, $|\varepsilon_{k',R}|$ and the regime–specific marker (at the cap: the shadow value λ; inside the cap: the optimal pledge θ^*) – can identify c as a driver of robot acquisition. We can interpret this result as an application of Andrews et al. (2020b) who formalize "transparency" of a structural report as the extent to which a reader, possibly holding different assumptions than the researcher, can reconstruct the structural quantity of interest from the reported statistics. They, together with Andrews et al. (2020a), show how a small set of descriptive or auxiliary statistics can make structural estimates informative in this sense. Our diagnostic is an application: given our model, the triple $(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \lambda)$ at the cap (or $(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \theta^*)$ in the interior) uniquely pins down whether liquidity (c) rather than intangibles (h) is the main driver of cross-sectional differences in adoption responses. We formalize this in the next proposition,[16]

Proposition 3. *If both elasticities increase and the marker also increases, then this sign pattern is unique to a liquidity tightening ($c \downarrow$). Other single–primitive changes produce a different signature. Hence, the joint behavior of $(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \lambda)$ or $(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \theta^*)$ provides a simple, falsifiable diagnostic for liquidity versus intangibles.*

Thus, to examine the effect of primitives on the elasticities we can turn to the following regression

[16] A formal statement and proof are in Appendix §IA.3.12–§IA.3.13.

specification:

$$\ln k'_{it} \;=\; \alpha_i + \tau_t \;+\; \underbrace{\gamma_p \left(\ln p_{k,t} \times \text{Cash}_{i,t-1} \right)}_{p_{k,t} \text{ heterogeneity by cash}} + \underbrace{\gamma_R \left(\ln R_t \times \text{Cash}_{i,t-1} \right)}_{R_t \text{ heterogeneity by cash}} + X'_{it}\delta \;+\; \varepsilon_{it}.$$

Predictions from our model indicate that $\gamma_p > 0$ and $\gamma_R > 0$. The interpretation would be that cross-sectionally low-cash firms in our sample have larger-in-magnitude elasticities to both shocks.[17] Using the tables above we can derive analogous predictions by replacing Cash with estimates of λ, θ. By comparing the pattern of coefficient signs, we can reject the hypotheses that a specific primitive is the main driver of differences in response by firms to shocks to the price of robots or the rate R.

We do not observe λ directly, but we can use text–based measures of firms' own disclosures about liquidity and investment delays as proxies. In the empirical work below we therefore treat the Hoberg–Maksimovic debt-constraint measure, and its machine–learning extension by Linn and Weagley (2024), as reduced–form proxies for the firm–level shadow value of funds. Firms that explicitly state in their 10-K statements that they may have to delay or scale back investment due to limited financing are assigned a high "debt-constraint" score. Conceptually, these statements correspond a high λ in our model: the marginal value of relaxing the financing constraint is large because otherwise profitable projects must be postponed.[18]

3.7 Extensions to Multiple Periods

In our model cash is given exogenously. However, in response to shocks firms may alter their cash holdings to increase or decrease cash availability for future periods. In that case, there is the possibility of an additional cash-adjustment response to shocks, sometimes termed the Euler-channel. This

[17] Note that there are a couple of measurement differences with respect to the empirical section. First, we indicate "low-cash" firms through a *positive* "Cash Constrained" indicator. Second, our price instrument is the *Yen/USD* exchange rate, where a higher value is a lower Yen and implies *a lower robot price*. Together, this implies a positive coefficient on the price instrument still maps into a positive γ_p, but a *negative* coefficient on the interest rate instrument maps into a positive γ_R.

[18] Hoberg and Maksimovic (2015) exploit the Liquidity and Capital Resources subsection of 10-K MD&A, which SEC Regulation S-K requires to discuss challenges to liquidity, planned investment, and sources of financing.

possibility introduces another term in the expressions for elasticity. Thus, for a shock x to p_k or R:

$$\frac{d\ln i_0}{d\ln x} = \underbrace{-\frac{d\ln M}{d\ln x}}_{\text{invoice channel}} \underbrace{-\left(\frac{c_1}{c_0-c_1}\right)\frac{d\ln c_1}{d\ln x}}_{\text{Euler–saving channel}} .$$

where $i = 0$ is current investment and c_1 is the amount of saving brought into the next period. The Euler-saving channel term is non-zero if the shock x changes the firm's expectation that its financial constraint λ in the next period will differ from the λ in the current period, taking into account the effect of the shock on the current λ. If the price or rate shock is transitory—or permanent but shifts today and tomorrow equally so that $\frac{d(\lambda_t - \mathbb{E}_t[\lambda_1])}{d\ln c_t \cdot d\ln x} = 0$—the Euler-saving term is zero; only "news" that differentially changes expected tightness at $t+1$ activates it.[19]

4 Data and Summary Statistics

4.1 BoL Data on Automation

Our study introduces a novel method of measuring firm-level adoption of automation technologies using detailed import transaction data. Unlike prior U.S. research which has predominantly relied on aggregated industry statistics or limited survey data, we exploit the concentrated structure of global robot supply to trace adoption via customs records. The structure of the robot market aids identification: a small number of globally dominant suppliers,[20] HS codes that cleanly identify robot and automation equipment, and concentrated production in Japan create a tight mapping from yen–dollar movements to dollar prices for U.S. importers. This allows us to trace firm-level user-cost shocks for a specific, high-tech capital input far more precisely than for generic CAPX, where transaction-level prices and quality are typically unobserved.

Our primary data source is the S&P Global Panjiva database, which provides comprehensive transaction-level records of international shipments from 2007 (the first year data is available) through 2019. This database captures Bills-of-Lading (BoLs) from U.S. Customs and Border Protection

[19] Even then, reversal of the cross-sectional pattern would require implausibly large saving responses by high-cash firms.

[20] Six companies—ABB (Switzerland), Omron, Fanuc, Kawasaki, and Yaskawa (Japan), and KUKA (Germany)—collectively accounted for 44% of global robotics industry revenue as of 2017 (Bonfiglioli et al., 2024).

(CBP), covering over 70% of U.S. international trade activities by weight. Each transaction record includes: consignee (importer) details including names, addresses, and unique identifier, foreign shipper information, detailed goods description, import quantity, shipment arrival date, port information (lading and unlading) and product codes (6-digit HS code).[21]

The step-by-step procedure with the exact number of observations for each step (including the matching to Compustat data described below) is summarized in Table IA.1.3. From the raw import data, we classify industrial robots using HS code 847950 and adopt the broader automation technology classification methodology outlined in Acemoglu and Restrepo (2022). Appendix Table IA.1.4 provides a comprehensive list of HS codes associated with automation technologies. Firms are designated as "adopters" if they report at least one import transaction that contains any of these HS6 codes. To ensure accurate measurement of firm-level adoption, we exclude importing firms that are likely logistic intermediaries. This includes firms identified on the list of the largest logistics companies in the United States, compiled by Armstrong & Associates, Inc., a prominent third-party logistics market research firm. Additionally, we exclude firms whose importer names include terms such as "logistic", "distribution", or "freight". We also remove companies identified as robot integrators, i.e., firms that specialize in installing robots for clients. This is based on a comprehensive list of U.S. robot integration firms, created through an extensive search across multiple robotics-related websites.[22] After applying these filters, we have 463,609 transactions from 2007 to 2019, where a U.S. firm imports an HS6 code identified as automation technology. We collapse transactions within each year to ensure that adoption measures capture any transaction occurring annually.

Before, matching to Compustat firms, we validate our methodology using all automation transactions. We compare our transaction-level import data with the International Federation of Robotics (IFR) statistics, which are widely used in prior research (e.g. Graetz and Michaels (2018); Acemoglu and Restrepo (2020); Dauth et al. (2021)). IFR compiles data on robot installations through annual

[21]The BoL data is described in detail in Flaaen et al. (2023) and has been widely used in academic research (e.g., Ganapati et al. (2024), Ayyagari et al. (2024)).

[22]The list includes well-known names such as Fanuc Corp, KUKA Robotics, and Kawasaki Robotics, many of which are subsidiaries of major global robot manufacturers. If these firms are listed as the *importers*, we exclude them from our analysis, as our focus is on measuring firm-level adoption for the purpose of producing non-robotic goods. Thus this exclusion ensures that the analysis targets end-user firms adopting automation technologies, rather than firms involved in the robotics supply chain.

surveys of industrial robot suppliers worldwide,[23] collecting information on robot installations by country, application, and customer industry. This data is sourced directly from manufacturers via questionnaires and supplemented with secondary data from national robotics associations.

Since IFR data exclusively covers industrial robots rather than broader automation technologies, Figure 3 compares IFR data with both our shipment data for industrial robots (right panel) and our broader automation technology measure (left panel). The figure reveals strong correlations between our measures and IFR data from 2007 to 2019. Specifically, the correlation between our volume-based measure of robot imports and IFR's robot adoption data is 0.87 (right panel). The left panel shows an even stronger relationship (0.95 correlation) between IFR data and our broader automation technology measure. These strong and robust correlations affirm the validity of our novel measurement approach.[24]

While survey data like IFR provide a direct measure of adoption, they are often constrained by the scope of their surveys and the frequency of data collection. In contrast, our data offers a more comprehensive view, capturing actual transactions with consistent tracking over the 2007-2019 period. Additionally, our dataset includes detailed information on the timing, quantity, and source of automation technology adoption.

FIGURE 3 ABOUT HERE

Sample Construction: Panjiva-Compustat Link We are interested in understanding how the characteristics of firms that adopt automation differs from those that don't and hence we need to match the importers (consignees) to publicly listed companies. To do this, we rely on the unique company identifier (*conpanjivaid*), which can be mapped to *GVKEY* identifiers in the Compustat database using the S&P Global Market Intelligence Business Entity Cross Reference Service (BECRS).[25] To construct our final sample, we first aggregate the shipments to the parent

[23] IFR uses the International Organization for Standardization (ISO) definition of industrial robots, which defines them as "automatically controlled, reprogrammable, multipurpose manipulators, programmable in three or more axes, which can be either fixed in place or mounted on a mobile platform for use in automation applications in an industrial environment."

[24] Similar relationships are also observed when using the number of shipments instead of volume.

[25] BECRS provides crosswalks across different firm identifiers including links between Capital IQ identifier (*companyid*) and Compustat identifier (*GVKEY*). See Table IA.1.3 for matching success.

firm-year level. We keep only firms that can be linked to a Compustat GVKEY and drop firms from the finance sector (SIC 60-67), utilities sector (SIC 40-49), and public administration (SIC 90-99). There are about 53,000 total firm-year observations in our Compusat sample, with 5,442 firm-year cases where an import of automation technology is observed.

4.2 Descriptive Patterns

We first examine the geographic distribution of robot adoption to identify regions with the highest concentration of automation technology imports. Using the universe of transactions from 2007 to 2019, each with an associated import address, we analyze adoption dynamics across U.S. counties. Figure 5 presents heat maps depicting the total volume of imported automation technologies (left panel) and industrial robots (right panel) by county during this period. The maps reveal striking geographic clustering, with the darkest concentrations appearing in major metropolitan areas and manufacturing centers. Notable "hotspots" emerge around counties containing cities such as Los Angeles, Chicago, Detroit, and Miami-Dade, reflecting both major ports of entry and established manufacturing hubs. The right panel, focusing specifically on industrial robots, shows a similar but slightly more concentrated pattern, suggesting that robot adoption is even more geographically clustered than broader automation technology imports.

The map highlights the highly uneven distribution of automation adoption across the United States. Large swaths of rural and less industrialized counties show minimal automation imports (represented by the lightest colors), while a relatively small number of counties account for the vast majority of import activity. These findings align with those of Brynjolfsson et al. (2023), who use establishment-level survey data to document the existence of *robot hubs*, or the highly concentrated distribution of robots across CBSAs in the U.S. A similar geographic pattern emerges when using the number of shipments instead of volume, confirming that the concentration is not merely driven by a few large-volume transactions but reflects genuine differences in automation adoption intensity across regions.[26]

[26] If we match firms to their Core-Based Statistical Area (CBSA) locations, we find that the top decile of CBSAs accounts for 85% of the total volume of robot imports in the U.S., while the next two deciles account for 9% and 3%, respectively. Using shipment counts instead of volume yields similar results: 83%, 9%, and 3% for the top three deciles respectively.

FIGURE 5 ABOUT HERE

The geographic concentration of robot adoption is mirrored by concentration on the supply side. Our data's richness, offering detailed information on both shipping and importing firms including their names and countries of origin, allows us to examine the structure of global robot supply chains. Panel A of Table 3 lists the countries from which most automation technologies shipped to the U.S. originate, while Panel B focuses specifically on industrial robots. Both panels reveal strong concentration of imports from a small set of countries, with five countries accounting for approximately 85% of the volume and 80% of the shipments of imported robots. Japan emerges as the dominant supplier, accounting for 58% of total robot import volume, followed by Germany at 9%. This Japanese dominance reflects the country's historical leadership in robotics technology and manufacturing.

TABLE 3 ABOUT HERE

This supply concentration reflects the dominance of a few key manufacturers in the global robotics market, with Japanese firms playing a particularly prominent role. Panel C of Table 3 shows that the top five companies - ABB (Switzerland), KUKA (Germany), Yaskawa (Japan), Fanuc (Japan), and Kawasaki (Japan) - account for nearly 45% of global robot production and serve as the largest suppliers to U.S. firms. Notably, three of these five leading suppliers are Japanese companies, underscoring Japan's central position in the global robotics industry. Our data reflects the country of shipment origin rather than corporate headquarters. For example, ABB, a Swedish-Swiss multinational, ships 92% of its U.S. robot imports from China, reflecting the company's strategic investments in Chinese manufacturing to supply both local and international markets.[27]

Having established the geographic and supply-side patterns using our complete dataset, we now focus on the subset of firms that can be matched to Compustat to examine industry-specific adoption patterns. Panel A of Table 4 shows the sectoral distribution of automation and robot adoption

[27] In 2022, ABB invested $150 million to open a fully automated robotics factory in China (Shanghai), its largest market. Currently, 90% of ABB robots sold in China are locally manufactured.

among publicly traded firms in our sample.

The data reveals that automation adoption is heavily concentrated in specific sectors. Manufacturing emerges as the dominant sector for automation adoption. Not only does it have the highest within-sector adoption rates (0.24% for automation technologies and 0.07% for industrial robots), but manufacturing firms also constitute the overwhelming majority of adopters economy-wide: 71.21% of all automation technology adopters and 78.04% of all industrial robot adopters come from the manufacturing sector. This concentration reflects manufacturing's natural affinity for production automation and the direct applicability of robotic technologies to manufacturing processes.

Given manufacturing's dominance in automation adoption, we further examine adoption patterns within manufacturing subsectors. The data reveals interesting heterogeneity within manufacturing, with some subsectors showing particularly high propensities for automation adoption. Metal products and machinery emerges as a leading subsector, with 48% of firms (within that subsector) adopting automation technologies and 8% adopting industrial robots specifically. This subsector accounts for 17.44% of all automation adopters within manufacturing and 20.99% of robot adopters within manufacturing. The high adoption rates likely reflect the subsector's production processes, which often involve repetitive metalworking operations well-suited to robotic automation, such as welding, cutting, and assembly tasks. Computers and electronics shows similarly high within-subsector adoption for automation technologies (25%) but lower robot adoption rates (4%). However, this subsector represents a substantial 20.62% of automation adopters and 20.8% of robot adopters within manufacturing. The high automation adoption reflects the precision requirements and clean-room environments common in electronics manufacturing, where automated systems can achieve the consistency and contamination control necessary for semiconductor and circuit board production. The automotive subsector demonstrates the highest robot adoption intensity, with 48% of firms adopting automation and notably 14% adopting industrial robots, the highest robot adoption rate among manufacturing subsectors. Automotive accounts for 8.54% of automation adopters but a disproportionately high 17.41% of robot adopters within manufacturing. This pattern reflects the automotive industry's pioneering role in industrial robotics, where robots have been extensively used for tasks like welding, painting, and assembly since the 1960s.

Panels B and C of Table 4 list the top ten U.S. companies adopting automation and industrial robots, respectively, ranked by total shipment volume from 2007 to 2019. These firms span diverse industries, including consumer robots, medical equipment, automobiles, solar panels, and construction materials.

TABLE 4 ABOUT HERE

4.3 Measures Used in the Analysis

We construct several measures to capture distinct dimensions of automation technology adoption. To measure the extensive margin, we define $Automate_t$ as a binary indicator for positive automation shipments in year t. For the intensive margin, we use the number of Bills of ladings, denoted $Shipments$, where each shipment reflects a unique cargo entry recorded in Panjiva, regardless of volume. Given the highly skewed distribution of shipments and the large mass at zero, we use the Inverse Hyperbolic Sine (IHS) transformation following Cohn et al. (2022), to create $Shipments$ (IHS). To account for firm size, we construct *Shipments/Assets*, which scales the number of shipments by lagged firm assets (in thousands of USD). Finally, we develop a stock measure, *Shipments-Stock (IHS)*, defined as the cumulative sum of annual automation shipments up to year t, also transformed using the IHS.

As an alternative measure of the intensive margin, we use volume in *Twenty-foot Equivalent Unit* (TEU), a standard metric of unit of cargo capacity, based on the volume of a 20-foot container. Again we take the IHS of Volume, $Volume(IHS)$, to accommodate the zeros.[28]

We construct the firm-level covariates from the Compustat database: *Log(Assets)* is the natural logarithm of total assets, which serves as our measure of firm size; *Log(Age)* is the natural logarithm of the number of years since the parent firm first appeared in Compustat; *Leverage* is the ratio of total debt (the sum of debt in current liabilities and long-term debt) to total assets; and return on assets *(ROA)*, calculated as net income divided by total assets.

[28] Each BoL represents a shipment, potentially containing multiple products (at the HS6 level). A shipment is classified as automation-related if any listed HS6 code corresponds to an automation good. However, this multi-product case accounts for only about 10% of automation imports, and often involves related HS codes. To compute TEU volume for each product (e.g., robots), we allocate container volume proportionally. For *Shipments*, we count one shipment per product. Around 5% of shipments lack TEU data, typically because they are not containerized. The correlation between *Shipments* and *Volume* is 0.98 over the sample period.

To capture liquidity constraints, we use the ratio of cash and short-term investments to total assets $Cash/Assets$. Firms with low cash holdings relative to assets are more likely to face internal liquidity constraints, limiting their capacity to finance investment projects without external capital. Liquidity-constrained firms may be particularly vulnerable to macroeconomic shocks, as they are less able to buffer investment during periods of heightened uncertainty or restricted credit availability.

To measure financing constraints, we employ Linn and Weagley (2024) debt constraints as our primary measure. The authors use a random forest model to predict the text-based constraints by Hoberg and Maksimovic (2015) using only observable accounting variables from Compustat, which allow them to extend and update up to 2019 the latter measures.[29] For periods for which both measures are avaialable, we show that our results hold using Hoberg and Maksimovic (2015) original measures.

For our empirical analysis, we construct binary indicators based on firm characteristics measured at the time of initial sample entry: *High Cash Constr* equals one if a firm's cash-to-assets ratio is below the sample median, indicating limited internal liquidity; *High FC* equals one if a firm's Linn-Weagley predicted financing constraints exceeds the median; and *High FC (HM)* equals one based on the Hoberg-Maksimovic constraint measure. *High Leverage* equals one for firms with above-median debt-to-assets ratios. Using baseline (first-year) values ensures that these indicators are predetermined relative to firms' subsequent automation decisions, thereby mitigating potential endogeneity concerns that could arise if constraint measures varied with investment outcomes.

4.4 Exchange Rate Shocks

As reported in Table 3, Japan, the dominant global producer of industrial robots,[30] is the primary source of U.S. robot imports. This market concentration creates a unique opportunity to identify exogenous price shocks to automation technology through exchange rate fluctuations.

In November 2012, following the election of Prime Minister Shinzo Abe, the Japanese Yen began a sharp and sustained depreciation against the U.S. dollar. Abe's platform—centered on aggressive

[29] Their analysis include variables like cash flow ratios, cash-to-capital, dividends, Tobin's q, leverage, sales growth, age, and size—thereby capturing non-linear relationships that traditional linear models might miss.

[30] According with this report from IFR delivered 45% of the global supply in 2022.

monetary easing to combat deflation (commonly referred to as "Abenomics")—marked a significant break from previous policy, triggering exchange rate movements largely unrelated to U.S. macroeconomic conditions or firm-level factors. Figure 4 plots the Yen–USD exchange rate. Between 2012 and 2015, the Yen depreciated approximately 33%, substantially reducing the dollar cost of Japanese robots.

FIGURE 4 ABOUT HERE

This exchange rate shock provides a useful natural experiment for three reasons. First, the concentration of robot production in Japan implies that Yen movements translate directly into price changes for a large share of automation technologies available to U.S. firms. Second, the timing and magnitude of the depreciation were driven primarily by Japanese monetary policy rather than U.S. economic fundamentals, supplying plausibly exogenous variation relative to U.S. investment decisions. Third, the shock was persistent, unfolding over several years and giving firms sufficient time to adjust automation strategies.

Although robot prices declined generally during this period due to technological progress, the post-2012 Yen depreciation generated an additional, exogenous reduction in the dollar price of Japanese robots. Given Japan's dominant market position, this depreciation amounted to a broad-based price shock in the U.S. automation market. The magnitude of the shock—approximately a 33% currency depreciation—was substantial enough to affect firms' relative cost-benefit calculations regarding automation adoption.

To verify that our results are not driven by general dollar strength, we include the *BroadUSD* index from FRED as a control in our robustness analysis. This trade-weighted index serves as a placebo: if our findings merely reflect broad dollar appreciation, similar effects should appear when using *BroadUSD*. We divide both exchange rates, *YEN–USD* and *BroadUSD* by 100 (the Broad Index uses a base of 100) and then take logarithms. As additional robustness checks, we also examine the *RMB–USD* and *EME–USD* series—the Chinese Yuan–USD exchange rate and the Nominal Emerging Market Economies Dollar Index, respectively—both obtained from FRED.

4.5 Monetary Policy Shocks

Automation investment entail large, long-lived capital outlays, making them particularly sensitive to fluctuations in the cost of capital. Assessing how interest rates affect such adoption requires isolating exogenous variation in borrowing costs. This is especially important because monetary policy is typically endogenous to broader macroeconomic conditions that affect firm investments.

We address this identification challenge using a high-frequency instrumental variables strategy. Specifically, we exploit variation in the 1-year Treasury yield around FOMC announcements and instrument it with cumulative monetary policy surprises following Bauer and Swanson (2023).[31] This strategy builds on a growing literature (e.g., Gertler and Karadi (2015), Ottonello and Winberry (2020), Döttling and Ratnovski (2023)) that leverages the near-instantaneous response of financial markets to policy news to ensure quasi-experimental variation.

Surprises are constructed using intraday changes in interest rates within a 30-minute window around FOMC announcements, as in Swanson and Jayawickrema (2024). Because policy decisions are finalized hours before their public release, this window cannot reflect policymakers' reactions to contemporaneous market conditions, mitigating reverse causality concerns. Compared to standard measures, the Bauer and Swanson (2023) surprises offer two key advantages. First, they expand the set of monetary policy signals by incorporating all public communications from the Fed Chair, substantially increasing the number and informativeness of events. Second, they orthogonalize raw surprises with respect to pre-announcement macro and financial variables, addressing the "Fed information effect" and ensuring that shocks capture unanticipated policy changes rather than systematic responses to incoming economic news.

Following Gertler and Karadi (2015), we cumulate these orthogonalized surprises to account for the gradual transmission of policy to financial conditions. The instrument is lagged with respect to the 1-year Treasury yield to ensure we identify the effect of policy on subsequent investment rather than contemporaneous reactions. While our regressions use unscaled yields, impulse responses are normalized to a 25 basis point change, consistent with standard practice (e.g., Döttling and

[31] We use the orthogonalized surprises series from the FRBSF Monetary Surprises Data.

Ratnovski, 2023).

This approach addresses key concerns in monetary policy identification. High-frequency surprises are orthogonal to both firm-level and macroeconomic confounders. Orthogonalization removes spurious correlations stemming by the Fed's endogenous response to information. Finally, the 1-year yield captures the relevant financing horizon for automation investments, which typically entails front-loaded capital expenditures and multi-year payoffs.

Data on Treasury yields and macroeconomic controls—including CPI inflation, industrial production, and real GDP growth—are sourced from the Federal Reserve Economic Data (FRED).

5 Main Results

This section presents our core empirical findings on firm-level automation adoption. We begin by establishing which firm characteristics systematically predict automation adoption using cross-sectional variation. We then examine how these patterns change in response to two major macroeconomic shocks: the Japanese Yen depreciation that dramatically reduced robot prices, and monetary policy surprises that altered the cost of capital for investment. Our results reveal striking heterogeneity in how firms respond to these shocks based on their financial constraints.

5.1 What types of firms Automate?

In this subsection, we investigate which firm characteristics are systematically associated with automation technology adoption decisions. We begin with a baseline regression specification which regresses various measures of automation adoption on firm characteristics while controlling for industry-by-year fixed effects to isolate cross-sectional variation in adoption patterns:

$$\mathrm{Y}_{it} = \beta_0 + \beta_1 \boldsymbol{X}_{i,t-1} + \alpha_{jt} + \epsilon_{it}. \tag{5.1}$$

where Y_{it} captures automation adoption by firm i in year t. We use five complementary proxies for automation adoption to ensure robustness: a binary adoption indicator(*Automate Dummy*), IHS-transformed shipment counts (*Shipments (IHS)*), shipments of automation technology per thousand

dollars of assets (*Automation Intensity*), cumulative automation stock (*Shipments-Stock (IHS*), and imported automation volume, (*Volume (IHS)*), the latter two measures also IHS-transformed. The firm characteristic vector, $X_{i,t-1}$ includes fundamental attributes such as size (*Log(Assets)*), age (*Log(Age)*), financial leverage (*Leverage*), and profitability (*ROA*), all measured with a one-year lag to mitigate endogeneity concerns. α_{jt} represents industry (NAICS3) x year fixed effects to absorb sector-specific trends and shocks.

The results are reported in Table 5. Panel A presents results for the full sample of firms across all industries. The findings reveal several consistent patterns across all five measures of automation adoption: Larger firms are significantly more likely to adopt automation technologies. The coefficients are positive and highly significant across all specifications, all statistically significant at the 1% level and economically significant.[32] Furthmore, it is reassuring that the several intensive margin effects tell a similar quantitative story, and one consistent with the extensive margin. This size effect becomes even more pronounced in the manufacturing subsample (Panel B). This suggests that firm scale plays an important role in automation decisions, likely due to larger firms having greater resources to invest in automation technologies and potentially greater ability to achieve economies of scale from such investments. This pattern aligns with prior European evidence from Koch et al. (2021) and Bonfiglioli et al. (2024), as well as U.S. studies by Brynjolfsson et al. (2023) and Firooz et al. (2025) establishing firm size as a consistent predictor of automation adoption across different contexts and methodologies.

TABLE 5 ABOUT HERE

Table 5 also shows that older, more established firms show higher rates of automation adoption, both across the full sample and the manufacturing sector. This could reflect older firms having more established processes that benefit from automation, greater accumulated capital, or more sophisticated management capabilities. .

We also see that firms with higher leverage ratios are more likely to adopt automation technologies.

[32]For example, a 1 SD increase in log assets is associated with almost a one-third standard deviation increase in the probability of importing automation technology. The quantitative magnitudes are very similar for the intensive margin.

The positive coefficients across all specifications suggest that leveraged firms may use automation as a strategy to improve operational efficiency and service their debt obligations more effectively. Tis finding suggests that automation serves not only as a growth strategy but also as a survival mechanism for leveraged firms seeking to enhance productivity and cash flow generation.

Interestingly, the relationship between profitability and automation adoption is negative and significant across all measures, indicating that less profitable firms are actually more likely to adopt automation. While this appears to contradict our theoretical prediction that more productive firms should automate more, it likely reflects heterogeneous motivations for automation that extend beyond our model's productivity-focused framework.[33]

Overall, the results show that automation adoption is primarily driven by firm scale, maturity, and financial constraints rather than superior performance. The findings suggest that automation may serve as both a growth strategy for large, established firms and a survival mechanism for less profitable companies seeking to improve their competitive position. The robustness across the five different automation measures indicates that these relationships reflect genuine economic patterns rather than measurement artifacts.

Given that manufacturing accounts for over 70% of automation adoption in our data and represents the primary use case for industrial robotics, we focus our subsequent macroeconomic shock analysis exclusively on manufacturing firms to ensure a homogeneous sample where automation technologies have clear and comparable applications across firms.

5.2 Automation Response to Macroeconomic Shocks

We now examine how firms' automation decisions respond to macroeconomic shocks. As seen in Figure 3, both the IFR data and our import data suggest a large increase in automation adoption between 2011 and 2018. In this section, we focus on two distinct but potentially complementary

[33]The negative relationship between profitability and automation adoption may reflect several factors: (1) heterogeneous motivations, where less profitable firms pursue automation as a survival strategy to reduce costs and improve efficiency, while profitable firms face less pressure for disruptive changes; (2) lifecycle effects, where mature firms with declining profitability automate to maintain competitiveness; (3) temporal mismatches, where current profitability may not capture forward-looking productivity considerations driving automation decisions; and (4) automation serving dual roles as both a growth strategy for established firms and a survival mechanism for struggling companies seeking operational transformation.

channels to explain this variation: capital price shocks that directly affect the cost of acquiring automation technologies, and interest rate shocks that influence the cost of financing such investments.

The key insight motivating this analysis is that firm responses to these shocks need not be homogeneous. As established in the corporate finance literature since (Bernanke et al., 1999), financial constraints fundamentally shape how firms respond to macroeconomic changes. Our theoretical framework predicts that there are financially constrained "marginal adopters", which are looking to scale up their technology adoption when costs decrease and their constraints are relaxed. These firms will be more sensitive to macroeconomic conditions when making automation decisions. To test this hypothesis formally, we estimate variants of the following specification:

$$\mathrm{Y}_{it} = \beta_0 + \beta_1 \boldsymbol{F_i} * \mathbf{Shock}_t + X_{it} + \alpha_i + \alpha_t + \epsilon_{it}. \tag{5.2}$$

where F_i represents a time-invariant measure of firm i's financial constraints, $Shock_t$ captures our macroeconomic variables of interest, and we include fixed effects for both firm and year to control for unobserved heterogeneity and aggregate trends. We also include the baseline firm controls (size, age, profitability, and leverage) in X_{it}.

5.2.1 Automation and Exchange Rate Shocks

Our first natural experiment exploits the sharp depreciation of the Japanese Yen between 2012 and 2015. Given Japan's dominance as a robot supplier (accounting for over 57% of U.S. industrial robot imports), Yen movements provide direct variation in robot prices that is plausibly exogenous to U.S. firm characteristics. This will allow us to test directly the predictions in Propositions that pertain to the price elasticity.

Panel A of Table 6 presents results where we replace year fixed effects with the YEN-USD exchange rate and macroeconomic controls (inflation rate, industrial production index, and real GDP growth) to directly estimate the price effect of currency movements. The results provide strong evidence that firms systematically respond to robot price changes driven by exchange rate movements. Across all three measures of automation adoption, the coefficients on YEN-USD are positive and highly

significant. The economic magnitudes are substantial: a 10% depreciation of the Yen against the dollar (corresponding to a 0.1 increase in our logged exchange rate measure) increases the probability of automation adoption by 2.25 percentage points (column 1), raises the inverse hyperbolic sine of automation shipments by roughly 4.2% (0.0424 IHS units, column 2)[34], and increases shipments relative to firm assets by 0.0525 shipments per thousand dollars of assets (column 3). For all 3 outcomes, applying the full 30% depreciation of the Yen implies an effect around 40-60% of the mean outcomes, especially important given automation's discrete, infrequent nature. These aggregate effects set the stage for examining whether such responsiveness varies systematically across firms based on their financial constraints.

TABLE 6 ABOUT HERE

Panel B examines whether the exchange rate effects documented in Panel A vary systematically between firms with different constraint profiles. By including both firm and year fixed effects, we absorb all aggregate trends and focus exclusively on the differential responses of high-constrained versus low-constrained firms to the Yen depreciation. The interaction terms reveal striking heterogeneity in firms' responsiveness to the price shock. High-cash constrained firms, i.e. those above the median cash-to-assets ratio in their first year of sample appearance, exhibit significantly stronger automation responses to Yen depreciation than their low-constrain counterparts. The interaction coefficients are positive and economically meaningful across all measures. In the last 3 rows we repeat the same specification for the constraint measures from Linn and Weagley (2024), which predicts constraints using observable accounting variables. We find very similar effects – firms identified as highly debt-constrained show significantly greater automation responses to Yen depreciation across all outcome measures and with comparable magnitudes – albeit in a smaller sample.

To ensure these effects reflect Japan-specific price changes rather than broader dollar strength, we control for the interaction with other exchange rates as well in Appendix Table IA.2.1. Specifically, we repeat the baseline specification while including interactions with the Chinese Renmibi, an index for emerging economy exchange rates, and the broad USD index. This test yields reassuring results:

[34]This can be interpreted as a log approximation for the positive observations, but obviously our analysis contains many zeros. Still, we note that using a $ln(1+x)$ transformation results in remarkably similar coefficients.

the differential effects persist when controlling for movements in these other currencies, with the high-constraint interaction with the YEN - USD exchange rate term remaining statistically significant and economically similar in magnitude across specifications.

The magnitude of these differential effects estimated in Table 6 Panel B is striking. For context, the interaction coefficient of 0.161 in column 1 implies that a 10% Yen depreciation increases the probability of automation adoption by an additional 1.61 percentage points for highly constrained firms relative to unconstrained firms, a substantial effect given that the baseline automation adoption rate in our sample is relatively low. The 30% depreciation experiences from 2012-2015 implies a total additional impact on constrained firms of 4.8 percentage points, about 1/3 of the mean adoption rate. The effects for the intensive margin outcomes are very similar in economic magnitudes – also at 30-50% of the mean number of shipments and shipments per thousand dollars of assets respectively – confirming that the differential sensitivity of constrained firms is robust across multiple measures of adoption.

As an alternative quantification, we decompose the average effect from Panel A. Constrained firms are approximately *twice* as responsive to Yen depreciation as unconstrained firms across adoption probability, shipment intensity (1.9 times), and asset-scaled automation (2.7 times).[35]

Event Study To trace out the evolution of the differential response between constrained and unconstrained firms, we estimate the following event-study specification:

$$Y_{it} \;=\; \sum_{\substack{\ell=-4 \\ \ell\neq 0}}^{8} \beta_\ell \left[\mathbf{1}\{\text{year}_t = 2011+\ell\} \times \text{FC}_i\right] \;+\; \mathbf{X}'_{i,t-1}\theta \;+\; \alpha_i \;+\; \varepsilon_{it}. \tag{5.3}$$

Where Y_{it} represents the automation outcome for firm i in year t and FC_i is a binary indicator equal to one if firm i's level of financing constraints in the first year it enters the panel was above the median for that year. The coefficients β_ℓ trace the differential path for financially constrained firms

[35]The calculation uses the fact that 40% of the regression sample is constrained firms (we use the median across all firms in our Compustat sample). Thus, using the 0.225 average effect (Panel A) and the 0.161 differential effect (Panel B), the total effect from a Yen depreciation for constrained firms is 0.322 and for unconstrained firms it is .16. The same calculation for debt constraints yields differential responses from 1.6 to 2.1 times larger for constrained firms.

relative to unconstrained firms in event-time ℓ years from the 2011 baseline (with $\ell = 0$ excluded as the reference period). We include lagged firm controls ($X_{i,t-1}$: log age, log assets, leverage, and ROA) and firm fixed effects (α_i) with standard errors clustered at the firm level.

This specification allows us to examine both the timing and persistence of differential automation responses while testing for pre-existing trends that might confound our interpretation. The choice of 2011 as the baseline year is motivated by the timing of Japan's monetary policy shift: Prime Minister Shinzo Abe's election in late 2012 and subsequent implementation of "Abenomics" created the sharp Yen depreciation that serves as our price shock.

Figure 6 presents the estimated event-study coefficients for the differential effects of automation adoption for financially constrained firms based on their cash liquidity levels.[36] The results provide compelling evidence for a causal interpretation of our exchange rate effects. Prior to 2012, the differential adoption patterns between high and low cash-constrained firms show no discernible trend, with confidence intervals that include zero and point estimates that fluctuate around zero without systematic direction. This absence of pre-trends is important for validating our identification strategy, as it suggests that financially constrained firms were not already on differential automation trajectories before the Yen depreciation began.

Starting in 2012, however, we observe a sharp and persistent divergence. The magnitude and persistence of these effects over roughly six years indicates that the Yen depreciation had lasting impacts on automation adoption patterns of firms with limited internal liquidity rather than merely shifting the timing of predetermined investments. Figure 6a suggests that the macroeconomic events after 2011 led to an increase in the number of shipments of automation technology imported that was 10-18% larger for constrained firms over the following eight years. Similarly, Figure 6b suggests that the large YEN depreciation associated with Japan's "Abenomics" policy increased the intensity of firm automation more so for cash constrained firms than for low constrained firms.

To further validate that our results reflect genuine responses to robot price changes rather than spurious correlation with other trends, Figure 7 presents equivalent event studies for general capital

[36] The results are qualitatively the same using other measures of firm financial constraints.

expenditure (CAPX) and research and development (R&D) spending. These placebo tests reveal markedly different patterns that strengthen confidence in our automation-specific findings. For both CAPX and R&D, we observe evidence of pre-existing trends that predate the 2012 Yen depreciation, suggesting these broader investment categories were already diverging between constrained and unconstrained firms for reasons unrelated to Japanese robot prices. Moreover, the post-2012 effects for these alternative outcomes are notably weaker and less persistent than those observed for automation.

FIGURES 6 AND 7 ABOUT HERE

This contrast reinforces our interpretation that the automation results reflect genuine responses to Japan-specific input price changes rather than broader differential investment trends between financially constrained and unconstrained firms. The clean event-study patterns for automation, combined with the noisier patterns for general investment, provide strong support for our identification strategy and highlight automation's particular sensitivity to input price shocks.

5.2.2 Addressing Selection

Selection on Constraints An important concern we take on is that the constraint indicator proxies for other firm attributes that also predict adoption responsiveness. We implement two standard selection-on-observables adjustments: inverse probability weighting (IPW) and propensity-score matching (PSM).

We first estimate each firm's likelihood of being classified as financially constrained via a logit model based on covariates such as size, age, profitability and Tobin's Q, controlling for industry and year effects. This is done on firm-level data from 2000-2007, thus outside of our baseline sample.[37] The weight for a constrained (treated) observation is $(1/\hat{p})$; for an unconstrained observation it is $1/(1-\hat{p})$. To limit variance inflation, we (i) inspect the weight distribution and (ii) verify post-weighting balance with standardized mean differences.[38]

Given the propensity scores, Panel A of Table 7 presents the results when re-estimating our baseline

[37] We can do this with Compustat data where robot adoption is not used. We match the average (pre-sample) firm score with the firm in the baseline sample, matching about three-quarters of firms. Thus the sample is reduced in this robustness check.

[38] In the balanced sample, the mean bias is 5%, below the 10% rule of thumb threshold.

panel specification with inverse propensity weights (IPW). This method preserves all observations but changes each observation's influence so the treated and control groups are comparable on observables in expectation. Re-estimating our baseline specification in this pseudo-population yields coefficients on the constraint–exchange rate interaction that remain positive and statistically significant.

As a complementary test, we match each constrained firm to its nearest unconstrained counterpart with a similar estimated propensity score. This matched sample restricts attention to firms that are observationally similar in size, profitability, and Tobin's Q in the 2000-2007 period. We restrict the analysis to pairs that lie on common support, ensuring direct like-with-like comparisons and run the same fixed-effects regression on the matched sample. Conceptually, this targets an ATT for constrained firms: among treated firms that have close analogs in the control group, how much more do they adopt in response to the macroeconomic shock? Panel B of Table 7 presents these results. The interaction coefficients again remain large and significant, indicating that the differential sensitivity of constrained firms is not an artifact of observable differences in fundamentals.[39]

Overall, the persistence of the constraint effects across both weighting and matching estimators suggests that heterogeneity in adoption reflects binding financial frictions rather than compositional differences across firms.

Selection on the Adoption Margin A related concern is that the heterogeneous responses we document may simply reflect two distinct regimes: firms that eventually adopt automation and those that never do, rather than within-firm differences in financing conditions. If non-adopters are structurally different (e.g., smaller, less export-oriented, in sectors with limited automation potential), then differences attributed to constraints could instead capture permanent adoption propensities.

To address this possibility, we control explicitly for firms' baseline probability of adoption and re-estimate our main specification within the subset of firms that lie on comparable adoption support. We first estimate a firm-level adoption score, defined as the predicted probability of ever adopting automation, using *pre-shock* (2007-2010) firm characteristics and industry–year effects. This score

[39]The results with the debt constraint indicator do become statistically insignificant in some cases, but this seems mostly due to a loss of power with a severely reduced sample size. The coefficients are consistently positive and economically large.

summarizes all observable factors that systematically distinguish adopters from non-adopters. We then implement two complementary tests reported in Appendix Table IA.2.2.

First, we re-estimate the specification restricting the sample to firms above the 25th percentile of the adoption-score distribution, that is, to firms that are potential adopters. This trims away firms whose ex-ante probability of adoption is effectively zero, isolating the intensive margin of investment responses among plausible adopters. The results are unchanged: the interaction between financial constraints and the exchange-rate shock continues to predict higher automation shipments and investment intensity. In the last 3 columns, we use a slightly different strategy to compare firms only to their closest counter-parts. With the full sample, we instead include industry times year fixed effects, and again the results are robust.

In the second panel (B), we include the adoption score directly in the baseline regression - both in levels and interacted with the exchange-rate shock - to absorb any component of the response attributable to inherent adoption likelihood rather than financing. Notice this is quite a restrictive test: we effectively control for differential responses taking into consideration an array of firm characteristics that predict adoption absent the shock (i.e. firm size, profitability, age). The coefficients on the constraint–exchange-rate interaction remain positive and significant – in fact of very similar magnitude – after this correction, implying that constrained firms respond more strongly even among those with similar underlying adoption propensities.

Taken together, these tests rule out a mechanical "two-regime" explanation. The stronger sensitivity of constrained firms is not driven by the extensive margin of adoption, whether a firm ever automates, but by differential investment responses within the set of firms that are already viable adopters.

5.2.3 Automation and Monetary Policy Shock

Our second natural experiment exploits high-frequency monetary policy surprises around Federal Reserve announcements, following Bauer and Swanson (2023). These surprises provide plausibly exogenous variation in interest rates that allows us to identify how the cost of capital affects automation decisions. Unlike the exchange rate shocks that directly affected input prices, monetary

policy changes influence the financing costs associated with large capital investments like automation technologies. As interest rates decline, the present value of future cash flows from automation investments increases, making robot adoption more economically viable across firms. However, this relationship is likely to be heterogeneous across the corporate landscape, with financing frictions playing a key role in determining which firms respond most strongly to interest rate fluctuations.[40]

Financially constrained companies face higher effective costs of capital during normal times, or liquidity shortages, creating a wedge between their investment thresholds and those of unconstrained firms. When interest rates decline, this wedge narrows, potentially triggering automation investments that were previously below the hurdle rate.[41] This mechanism aligns with the broader literature on financial accelerator effects (Gertler and Karadi, 2011), where financial frictions magnify the impact of monetary policy shocks on real economic activities. Our model makes specific predictions about the rate elasticities: that more constrained firms will respond more aggressively to financing shocks along with price shocks.

We follow the empirical strategy described in equation 5.2. This time, the shock are changes in the 1-year Treasury rate (lagged), which we instrument using the monetary policy surprises of Swanson and Jayawickrema (2024). In the baseline specification in Table 8, Panel A we see that interest rate increases systematically reduce automation adoption. Higher Treasury rates significantly decrease automation across all measures, with first-stage F-statistics exceeding 2,000, indicating strong instrumental relevance. However, these aggregate effects mask substantial heterogeneity across firms as seen in Panel B of Table 8.

TABLE 8 ABOUT HERE

The interaction specifications in columns (1)-(3) reveal that these negative effects are heavily concentrated among high-cash constrained firms, while low-constrained firms appear largely insulated from interest rate changes. The interaction coefficient of -0.085 for automation adoption indicates that a 25-basis-point increase in the 1-year Treasury rate reduces automation probability by an

[40] While Döttling and Ratnovski (2023) demonstrate that monetary easing disproportionately stimulates intangible investment, we extend this logic to automation technologies, which are closer to tangible investment.

[41] Unlike Fornaro and Wolf (2022), who model aggregate automation responses to monetary policy without firm-level heterogeneity, our analysis highlights how financial constraints amplify the automation-interest rate sensitivity.

additional 2.1 percentage points for high-constrained firms relative to low-constrained firms. Once again, constrained firms respond at about 2.5 larger rates than unconstrained firms. Similarly, the effects on automation intensity and asset-scaled automation show that constrained firms exhibit substantially stronger negative responses to interest rate increases.

The next three columns replicate the analysis using our broader measure of debt constraints from Linn and Weagley (2024). Highly debt-constrained firms show significantly larger negative responses to interest rate increases across all automation measures, though the asset-scaled measure becomes statistically insignificant while maintaining the expected negative sign. The consistency between our constraint-based results reinforces our interpretation that financial limitations systematically amplify firms' sensitivity to monetary policy changes.

Impulse Responses To examine the persistence and timing of monetary policy effects, we employ instrumental-variable local projections following Jordà (2005). For each horizon h, we estimate:

$$y_{it+h} - y_{it-1} = \beta_1^h \cdot \hat{R}_t + \gamma_1^{h'} \cdot X_{t-1}^m + \gamma_2^{h'} \cdot X_{it-1} + \mu_i + \varepsilon_{it} \tag{5.4}$$

where the outcome variable y_{it} is a measure of the investment of firm i in year t in automation. $\hat{R}_t$ is the negative of the instrumented 1-year Treasury rate. X_{t-1}^m are lagged macroeconomic control variables (log industrial production, real GDP growth, and inflation rate), X_{it-1} are lagged firm-level controls including Log(Assets), Log(Age), Leverage, and Profitability, and μ_i are firm-fixed effects.

Panel (a) of Figure 8 presents the impulse response of automation investment (scaled by assets) to a 25 basis point decrease in the 1-year Treasury rate for all manufacturing firms. The results show that automation investment experiences a significant positive response that persists for multiple years. Specifically, automation investment increases by approximately 1.2 percentage points the year of the shock, remains significantly positive at around 1 percentage point in years 1-2, and gradually decreases remaining positive (though statistically insignificant) by year 4.

The aggregate effects mask substantial heterogeneity across firms with different financial profiles. We first split the sample into high- and low-leverage firms, based on each firm's cash-constrain levels

in the first year a firm appears in the data relative to the median firm in that year.

Panel (b) of Figure 8 reveals that the increase in automation response is concentrated more so among financially constrained firms. For these firms, automation investment increases by approximately 2 percentage points immediately following the falling rates and while decreases over time, remains significantly positive for the following three-year period, with a roughly 1 percentage point effect throughout. The persistence and magnitude of this response underscore how financing frictions amplify the transmission of monetary policy to automation decisions. In stark contrast, Panel (c) of the same Figure, shows that low-constraint firms exhibit a small and marginally significant response to interest rate changes.

FIGURES 8 ABOUT HERE

Finally, we compare the response to a monetary policy shock of firms' investments in automation technology with their response to other important investment decisions of firms. Figure 9, shows the impulse response function for CAPX and R&D spending, both for high- and low-constrained firms. Panels (a) and (b) show the results for CAPX, which considers firm investments in long-term assets, such as property, plant, and equipment (including automation machinery). Panel (a) shows that firms with high constraints increase their capital investment to assets ratio by less than 0.1 percentage points right after a 25 basis point reduction in the interest rate, but the effect completely dissipates for the following years. Since the units of automation are in Shipments and CAPX in value, the magnitudes are not comparable. However, we can say with certainty that the effect on CAPX is contemporaneous and transitory, whereas the effect on automation is persistent for at least four years. Panel (b) shows that firms with low constraints have no significant effect for the next four years after the shock.

Panels (c) and (d) show the response on R&D investment to the monetary policy shock for firms with high and low constraints, respectively. In both cases, the responses are null and negative for low constrained firms for the years two to four after the shock. While this response contrasts with firms' automation investment adjustments, it is consistent with the fact that firms in the first stages of the product life cycle need to invest in product development regardless of their financial constraints,

whereas firms that are ready to scale up production with automation technologies can always delay that important step and wait for better interest rates.

FIGURE 9 ABOUT HERE

Alternative Measures of Shocks to Financing As an additional robustness check on our monetary policy results, Table IA.2.3 examines whether financially constrained firms' automation responses vary with alternative measures of credit market conditions beyond the Treasury yield. Panel A uses the Moody's Baa-10y spread—the difference between investment-grade corporate bond yields and 10-year Treasury yields—which captures the risk premium firms face when accessing debt markets. Panel B employs the Excess Bond Premium (EBP) of Gilchrist and Zakrajšek (2012), which isolates the non-default component of corporate bond spreads and serves as a purer measure of credit supply conditions and investor risk appetite. Consistent with our main findings, we find that tighter credit conditions disproportionately constrain automation investment among financially vulnerable firms. In Panel A, cash-constrained firms reduce automation adoption significantly more than unconstrained firms when credit spreads widen, with interaction coefficients ranging from -0.011 to -0.043 across the three outcomes. The effects are similarly negative for debt-constrained firms, though somewhat smaller in magnitude. Panel B shows that elevated excess bond premia likewise dampen automation more among constrained firms, with statistically significant negative interactions for both constraint measures across most specifications. These results reinforce our interpretation that financing frictions amplify the transmission of monetary conditions to automation investment, and demonstrate that our core findings are robust to using credit spread measures that reflect actual corporate borrowing costs rather than risk-free rates alone.

5.2.4 Responses to Both Shocks

Given the multiple concurrent shocks driving the automation adoption dynamics we observe, we next examine firms' responses when both exchange rate and monetary policy shocks occur simultaneously. As these shocks contain independent variation in terms of timing and magnitudes, we test whether both can explain part of the heterogeneous adoption responses documented in the previous sections.

This analysis is particularly important because firms in practice face both types of macroeconomic pressures simultaneously, and understanding their relative importance can inform both theoretical models and policy discussions. In fact, Tables 1 and 2 in the theory section show that the combination of price and rate elasticities allows us to speak to the specific mechanisms driving automation responses. For example, price and rate elasticities that move in the same direction suggest a liquidity effect and negate a story where initially more debt-constrained firms see their debt become relatively less risky (changes in y) or larger changes in their collateral cap ($\bar{\theta}$).

Table 9 presents results from specifications that include interactions of firm financial constraints with both the Yen-Dollar exchange rate and the instrumented Treasury rate. In all specifications, we include both firm and year fixed effects, controlling for any aggregate macroeconomic shocks, as we focus on the differential responses across firm types.

The results demonstrate that both macroeconomic factors play significant and economically meaningful roles in driving automation investment across firms based on cash constraints. For the exchange rate channel, high-constraint firms show consistently positive and significant responses to Yen depreciation across all automation measures, while for the monetary policy channel these firms exhibit significant negative responses to interest rate increases across most specifications. Broadly speaking, the interaction coefficients are in line with the results in Tables 6 and 8, with the relative price responsiveness now marginally larger and the relative interest rate responses slightly lower.

Again, we can quantify the relative responses of high vs. low constrained firms, and furthermore how this compares across the price and interest rate shock. From column (1) of Table 9, we use the two interaction coefficients as the differential effects, while normalizing the average effect of each shock with the results in Panel A of Tables 6 and 8 respectively. In this way, the adoption probability response of constrained firms to the price shock is 2.5 times larger than unconstrained firms. With respect to the monetary policy shock, the ratio is 1.6. This suggests that firms are more sensitive to equivalent exchange rate shocks than to monetary policy shocks, indicating that input price variation may be a more powerful driver of automation investment than financing costs. Across the 3 outcome measures, this conclusion is generally consistent, with the relative response to the

price shock being 45-60% larger.

5.2.5 Robustness

We subject the above regression results to several robustness tests. First, Table 10 provides important robustness evidence by replicating the analysis using alternative measures of financial constraints. The first 3 columns employ the Hoberg and Maksimovic (2015) text-based debt constraint measure, which captures firms' self-reported financing difficulties through systematic analysis of 10-K filings. These should be similar to Linn and Weagley (2024), but are only available until 2015.[42] Using this text-based constraint measure, we find qualitatively similar results to those in Table 9. In the last 3 columns we instead examine leverage ratios, defined as total debt relative to total assets of the firm. When we define firms as "high leverage" or "low leverage" using the familiar definition, the sign and magnitude of the adoption responses are once again very much in line with the baseline constraint measures.

TABLE 10 ABOUT HERE

The consistency of results across leverage ratios, text-based debt constraints, and cash constraints strengthens our interpretation that limited financial flexibility drives the enhanced responsiveness to macroeconomic shocks. Furthermore, Appendix Table IA.2.4 provides further robustness for the constraints using continuous constraint measures rather than binary indicators. Interacting the yen exchange rate and Treasury yield with cash-to-assets ratios (columns 1-3) and the continuous Linn and Weagley (2024) index (columns 4-6) yields consistent results: lower cash holdings and higher predicted constraints amplify firms' responses to both price and financing shocks. This demonstrates our findings are not sensitive to threshold choices in constructing constraint indicators. This robustness across different constraint measures suggests that our findings reflect genuine economic relationships rather than measurement artifacts from any particular constraint proxy.

Second, a potential concern with our main results is that automation adoption might be correlated with firms' broader globalization strategies, particularly their use of international outsourcing.

[42] We can still assign firms a fixed constraint measure if it is in the data both before and after 2015, but we cannot assign a constraint indicator for firms that enter the data after 2015.

Outsourcing represents an alternative strategy for firms to reduce production costs and improve efficiency. Rather than investing in automation technologies to replace labor, firms might choose to relocate production to lower-cost countries. This substitution possibility is particularly relevant for our analysis because both strategies respond to similar economic pressures - the need to reduce costs and improve competitiveness. Moreover, exchange rate movements that make automation cheaper might simultaneously affect the relative attractiveness of outsourcing to different countries, potentially creating spurious correlations in our results.

To address this concern, Table 11 presents robustness tests that control for firms' outsourcing activities. Following Hummels et al. (2014), we define outsourcing as imports in the NAICS industry that the firm operates in.[43] We construct two measures: an *Outsourcing dummy* indicator that equals one if the firm imports any products (excluding automation) in its industry in that year, and *Outsourcing intensity* measured as the number of import shipments in the firm's industry scaled by lagged assets.

TABLE 11 ABOUT HERE

In Panel A, examining heterogeneity by cash constraints, highly constrained firms continue to show significantly stronger responses to both Yen depreciation and Treasury rate changes when outsourcing controls are included. Panel B, using debt constraints, shows similarly robust results. The consistency of these results with our main findings indicates that outsourcing behavior does not drive our core conclusions about how financial constraints amplify automation responses to macroeconomic shocks.[44]

Third, a key question in the literature concerns which firms drive aggregate investment responses to macroeconomic shocks. Kroen et al. (2024) find that industry leaders benefit most from interest rate cuts, attributing this to duration effects on long-lived intangible assets such as brands and patents.

[43] This is constructed with the same import data as the baseline analysis, now capturing the broader set of goods firms import.

[44] Interestingly, the outsourcing interaction terms are themselves positive and significant, suggesting that the macro trends during this time also led firms that participate in international trade to adopt more automation technologies. This pattern is consistent with automation and outsourcing serving as complements rather than substitutes in firms' broader efficiency-enhancement strategies, possibly because globally engaged firms face more competitive pressures that drive adoption of multiple cost-reduction technologies.

To directly test whether our automation investment results concentrate among certain large and growing firms, possibly reflecting an alternative mechanism, we conduct several checks in Table 12.

In Panel A we exclude "star firms", defined as the top 5% of firms by market capitalization within each industry-year following Kroen et al. (2024) and examine responses among the remaining smaller manufacturers. The results reveal that financially constrained laggards, rather than industry leaders, drive automation adoption in response to macroeconomic shocks.[45] In Panel B, we add several controls to the baseline specification. Namely, we include Tobin's Q (Peters and Taylor (2017))[46]; the intensity of intangible capital usage—the ratio of intangible capital to total capital—to account for firms' knowledge-intensity and its potential correlation with both automation adoption and access to finance; and finally, controls for firms' sales growth quartile to ensure that our results are not driven by differential trends among high-growth firms. Across all specifications, our main results remain qualitatively and quantitatively unchanged, even with a significant drop in sample coverage. This comprehensive robustness table supports the paper's core argument that financially constrained firms exhibit heightened sensitivity to macroeconomic shocks due to their difficulty in raising non-pledgeable intangible capital, not due to diverging investment opportunities.

Table 12 About Here

The internet appendix includes several extra robustness tests that further reject alternative mechanisms. Most importantly, Appendix Table IA.2.5 conducts a horse race between our cash constraint measure and distance-to-default, an alternative metric of financial distress used in prior work (Gilchrist and Zakrajšek, 2012; Ottonello and Winberry, 2020). Distance-to-default exhibits no significant effects on automation responsiveness to either exchange rate or interest rate shocks, while cash constraints remain strongly significant, indicating that our results reflect liquidity constraints rather than default risk per se. Table IA.2.6 excludes automotive sector firms (NAICS 336), a particularly

[45] We don't view this as a conflict in mechanisms, but rather that Kroen et al. (2024)'s duration channel operates primarily among unconstrained firms with substantial intangibles, while our collateral channel operates among constrained firms investing in pledgeable automation equipment. The two mechanisms coexist but affect different segments of the firm distribution through different asset classes.

[46] Q theory posits that firms base their investment decisions on the ratio of market value to replacement cost of assets, and when $Q > 1$, firms should invest more since the market values their assets above replacement cost. Hence, it is essential to demonstrate that the effects of exchange rate and monetary policy shocks on automation investment persist even after conditioning on Q. We use *Total Q*, incorporating intangibles.

robot-intensive industry, and confirms that our findings are not driven by this specific subsector. Finally, Table IA.2.3 splits the sample into periods of declining interest rates (2007-2013) versus rate normalization (2014-2019). The negative interaction between firm constraints and monetary policy shocks is concentrated in the earlier period when rates were falling, consistent with the interpretation that automation responses are primarily driven by the easing of financing conditions.

6 Discussion

Large macro shocks that shift the price of imported automation capital can substantially reshape firms' adoption decisions. The 2012–2015 yen depreciation generated exactly such a shock: it sharply lowered the domestic price of Japanese robots and, in doing so, moved a wide set of firms closer to or past their internal adoption thresholds. Our estimates show that financially constrained firms are the most responsive to these price changes, suggesting that episodes of major exchange rate movements can accelerate the convergence of technologically lagging firms toward frontier production methods. More broadly, macroeconomic shocks can have persistent structural effects because they determine which firms cross their adoption thresholds, influencing both the level of aggregate investment and the distribution of technological upgrading across the firm distribution.

6.1 A Counterfactual Exercise: Automation Adoption Without the Yen Depreciation

To assess the economic significance of the heterogeneous elasticities we estimate, we implement a counterfactual simulation asking how many robots financially constrained firms would have adopted had the yen remained at its 2011 value. Holding the exchange rate constant eliminates the steep decline in robot prices that occurred between 2012 and 2015. Using the estimated coefficient on the interaction between *High Cash Constraint* and the yen, we recompute predicted robot imports for each firm-year. Financially constrained firms would have imported approximately 353 fewer robots in this counterfactual, a 17.3 percent shortfall relative to observed adoption while unconstrained firms exhibit essentially no change. This confirms that the price shock was central to the diffusion of automation among marginal adopters.

This aggregate effect masks substantial heterogeneity across industries as presented in Figure 10. In a number of price-sensitive sectors, the implied shortfall exceeds 20–30 percent (e.g. Chemical Manufacturing (NAICS 325), and Other Manufacturing (NAICS 339). Part of this variation reflects genuine differences in elasticity across types of production, and part reflects mechanical scale differences: industries with low baseline robot use exhibit large percentage swings even when absolute changes are modest. Taken together, the aggregate and industry-level results reinforce the conclusion that robot price movements rather than financing conditions were the dominant force driving automation adoption among financially constrained firms during this episode.

FIGURE 10 ABOUT HERE

A 17.3 percent decline is also large relative to what our estimates imply for financing channels. In Table 9, the High Cash Constraint interaction with the yen is roughly four to five times larger in absolute value than the corresponding interaction with the one-year U.S. Treasury yield. Scaling our 17.3 percent price effect by this ratio implies that even a sizeable 100 basis point monetary policy shock would move automation adoption among constrained firms by only about four percent, and the actual high-frequency policy surprises in our sample are smaller than this benchmark. Because robots are partially pledgeable and their prices are set in international markets, exchange-rate induced price movements enter the adoption decision more directly than modest changes in U.S. borrowing costs. This comparison highlights a central contribution of our paper: for financially constrained firms, the dominant margin is not interest rates but the price of imported automation capital. Exchange-rate movements therefore have first-order effects on the diffusion of frontier technologies, far exceeding the impact of contemporaneous monetary policy shocks.

6.2 Interpreting the Evidence Through the Model: The Roles of c and h

The empirical findings, including the counterfactual exercise, align closely with the mechanism articulated in our theoretical model. Proposition 3 provides a sharp diagnostic: only liquidity tightening (lower c) simultaneously raises both the price and rate elasticities and increases the relevant financial-tightness marker (λ at the cap or θ^* in the interior). No other primitive - including

a change in intangible costs h - produces this joint pattern. This is a testable prediction and our data strongly support it. Two implications follow.

First, the marginal adopters are low-c firms. These firms have high effective cash shares and are therefore most sensitive to changes in the sticker price net of collateralized borrowing. When robot prices fall, their internal "invoice" drops sharply, moving them across the adoption threshold. Our counterfactual confirms this mechanism: eliminating the yen depreciation reduces adoption among constrained firms by roughly 17 percent, exactly the pattern predicted for low-c marginal adopters.

Second, variation in h cannot rationalize the observed cross-sectional differences. Proposition 3 shows that lower intangible capital h would raise both elasticities while loosening the financial margin (lower λ or lower θ^*). Empirically, we see the opposite pattern: firms with the largest elasticities are precisely those that appear most financially tight. Together, these facts rule out an h-based mechanism and match the Proposition 3 signature of liquidity-driven heterogeneity.

In short, the joint behavior of (i) the price elasticity, (ii) the rate elasticity, and (iii) firm-level indicators of financial tightness line up exactly with the Proposition 3 signature of low-cash marginal adopters, not high-h firms. This tightly links the model to the data and provides a direct, falsifiable confirmation of the mechanism.

These findings contribute to several areas of research. They relate to the literature on heterogeneous monetary policy transmission. Unlike settings in which financially constrained firms respond less to positive shocks, our case shows the opposite. Because robots are partially pledgeable and priced in global markets, constrained firms respond more strongly to favorable price movements. This underscores the role of asset pledgeability and investment opportunities in determining the sign and magnitude of firm-level responses.

The results also relate to work on market structure and technological diffusion. If macro shocks disproportionately expand automation among constrained firms, they can compress technological gaps. If shocks raise financing costs or increase input prices, they can widen these gaps. Macroeconomic policies and external global shocks therefore meaningfully shape the trajectory of technological diffusion and market concentration.

Taken together, the evidence places our contribution at the intersection of financing constraints, collateral economics, and the diffusion of frontier production technologies. We show that price effects dominate financing effects, that constrained firms gained substantially from the yen depreciation, and that the empirical patterns are consistent with a marginal adopter who is low c rather than low h. This combination provides a unified explanation for why automation adoption remains uneven across firms and why macroeconomic shocks can either accelerate or slow the spread of transformative technologies.

7 Conclusion

This paper provides the first comprehensive evidence on how financial constraints condition firm responses to macroeconomic shocks in the adoption of transformative technologies. Contrary to the conventional view that constraints simply limit investment capacity, we show that constrained firms are systematically more responsive to both capital-price and financing shocks, making them the marginal adopters of automation when conditions shift.

Using novel transaction-level data on U.S. automation imports linked to Compustat, we exploit two natural experiments – Japan's yen depreciation (2012-2015) and high-frequency monetary policy surprises. Constrained firms are about two times more responsive to exchange rate and interest rate shocks than unconstrained firms. At the same time, exchange rate shocks exert relative effects that are 45-60% larger than rate shocks when normalized by the average effect of each shock, highlighting that variation in input prices dominates borrowing costs in shaping the investment dynamics we highlight in this paper.

We rationalize these patterns with a collateral model in which robots are pledgeable but entail non-pledgeable intangible capital. The model collapses adoption into an invoice identity, linking firm fundamentals directly to investment elasticities. Firms with higher effective cash shares on this invoice emerge as the financially marginal adopters, amplifying their sensitivity to both favorable price declines and looser borrowing conditions.

The broader implications are twofold. First, automation diffusion is disproportionately pulled

forward by constrained incumbents rather than unconstrained leaders, reshaping competitiveness and inequality. Second, the dominance of price over financing channels highlights the importance of examining multiple transmission mechanisms when assessing the effects of macroeconomic shocks. In an environment of ongoing trade tensions, targeted tariffs, and shifting monetary conditions, these results suggest that the collateral channel should be considered alongside traditional demand and discount-rate channels in both policy design and firm strategy.

References

Acemoglu, D., Anderson, G. W., Beede, D. N., Buffington, C., Childress, E. E., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., et al. (2022). Automation and the workforce: A firm-level view from the 2019 annual business survey. National Bureau of Economic Research.

Acemoglu, D., Lelarge, C., and Restrepo, P. (2020). Competing with robots: Firm-level evidence from france. *AEA Papers and Proceedings*, 110:383–88.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from us labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2022). Demographics and automation. *Review of Economic Studies*, 89:1–44.

Almeida, H. and Campello, M. (2007). Financial constraints, asset tangibility, and corporate investment. *The Review of Financial Studies*, 20(5):1429–1460.

Almeida, H., Campello, M., and Weisbach, M. S. (2004). The cash flow sensitivity of cash. *The journal of finance*, 59(4):1777–1804.

Amiti, M., Itskhoki, O., and Konings, J. (2014). Importers, exporters, and exchange rate disconnect. *American Economic Review*, 104(7):1942–1978.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2020a). On the informativeness of descriptive statistics for structural estimates. *Econometrica*, 88(6):2231–2258.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2020b). Transparency in structural research. *Journal of Business & Economic Statistics*, 38(4):711–722.

Antonioli, D., Marzucchi, A., Rentocchini, F., and Vannuccini, S. (2024). Robot adoption and product innovation. *Research Policy*, 53(6):105002.

Ayyagari, M., Cheng, Y., and Weinberger, A. (2024). Surviving pandemics: The role of spillovers. *Journal of Financial and Quantitative Analysis Forthcoming*.

Bauer, M. D. and Swanson, E. T. (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual*, 37(1):87–155.

Benmelech, E., Kumar, N., and Rajan, R. (2022). The secured credit premium and the issuance of secured debt. *Journal of Financial Economics*, 146(1):143–171.

Benmelech, E. and Zator, M. (2025). Robots and firms. *Journal of Financial Economics*, 159:112139.

Bernanke, B. S. and Gertler, M. (1995). Inside the black box: the credit channel of monetary policy transmission. *Journal of Economic perspectives*, 9(4):27–48.

Bernanke, B. S., Gertler, M., and Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics*, 1:1341–1393.

Bessen, J. E., Goos, M., Salomons, A., and Van den Berge, W. (2025). What happens to workers at firms that automate? *The Review of Economics and Statistics*, 107(1):125–141.

Blaum, J., Lelarge, C., and Peters, M. (2018). The gains from input trade with heterogeneous importers. *American Economic Journal: Macroeconomics*, 10(4):77–127.

Bolton, P., Chen, H., and Wang, N. (2011). A unified theory of tobin's q, corporate investment, financing, and risk management. *The Journal of Finance*, 66(5):1545–1578.

Bonfiglioli, A., Crino, R., Fadinger, H., and Gancia, G. (2024). Robot imports and firm-level outcomes. *The Economic Journal*, 134:3428–3444.

Brynjolfsson, E., Buffington, C., Goldschlag, N., Li, J. F., Miranda, J., and Seamans, R. (2023). Robot hubs: The skewed distribution of robots in us manufacturing. *AEA P & P*, pages 215–218.

Catherine, S., Chaney, T., Huang, Z., Sraer, D., and Thesmar, D. (2022). Quantifying reduced-form evidence on collateral constraints. *The Journal of Finance*, 77(4):2143–2181.

Chaney, T., Sraer, D., and Thesmar, D. (2012). The collateral channel: How real estate shocks affect corporate investment. *American Economic Review*, 102(6):2381–2409.

Cloyne, J., Ferreira, C., Froemel, M., and Surico, P. (2023). Monetary policy, corporate finance, and investment. *Journal of the European Economic Association*, 21(6):2586–2634.

Cohn, J., Liu, Z., and Wardlaw, M. (2022). Count (and count-like) data in finance. *Journal of Financial Economics*, pages 529–551.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Di Giovanni, J., Kalemli-Özcan, Silva, A., and Yildirim, M. A. (2022). Global supply chain pressures, international trade, and inflation. National Bureau of Economic Research.

Dixon, J., Hong, B., and Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605.

Döttling, R. and Ratnovski, L. (2023). Monetary policy and intangible investment. *Journal of Monetary Economics*, 134:53–72.

Drechsel, T. (2023). Earnings-based borrowing constraints and macroeconomic fluctuations. *American Economic Journal: Macroeconomics*, 15(2):1–34.

Edgerton, J. (2010). Investment incentives and corporate tax asymmetries. *Journal of Public*

Economics, 94(11-12):936–952.

Engineering.com (2018). Time and money: How much do industrial robots cost? . Accessed: Nov 20, 2025.

Falato, A., Kadyrzhanova, D., Sim, J., and Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the us corporate savings glut. *The Journal of Finance*, 77(5):2799–2852.

Farre-Mensa, J. and Ljungqvist, A. (2016). Do measures of financial constraints measure financial constraints? *The Review of Financial Studies*, 29(2):271–308.

Fazzari, S. M., Hubbard, R. G., and Petersen, B. C. (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity*, 19(1):141–206.

Firooz, H., Liu, Z., and Wang, Y. (2025). Automation and the rise of superstar firms. *Journal of Monetary Economics*, page 103733.

Flaaen, A., Haberkorn, F., Lewis, L., Monken, A., Pierce, J., Rhodes, R., and Yi, M. (2023). Bill of lading data in international trade research with an application to the covid-19 pandemic. *Review of International Economics*, 31(3):1146–1172.

Fornaro, L. and Wolf, M. (2022). Monetary policy in the age of automation. Universitat Pompeu Fabra, Department of Economics and Business, Working Paper.

Frank, M. Z. and Goyal, V. K. (2009). Capital structure decisions: Which factors are reliably important? *Journal of Financial Economics*, 93(1):1–23.

Ganapati, S., Wong, W. F., and Ziv, O. (2024). Entrepot: Hubs, scale, and trade costs. *American Economic Journal: Macroeconomics*, (4):239–78.

Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17–34.

Gertler, M. and Karadi, P. (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American economic review*, 102(4):1692–1720.

Gopinath, G., Kalemli-Özcan, Ş., Karabarbounis, L., and Villegas-Sanchez, C. (2017). Capital allocation and productivity in south europe. *The Quarterly Journal of Economics*, 132(4):1915–1967.

Gopinath, G. and Neiman, B. (2014). Trade adjustment and productivity in large crises. *American Economic Review*, 104(3):793–831.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

Hennessy, C., Levy, A., and Whited, T. (2007). Testing q theory with financing frictions. *Journal of Financial Economics*, 83:691–717.

Hoberg, G. and Maksimovic, V. (2015). Redefining financial constraints: A text-based analysis. *The Review of Financial Studies*, 28(5):1312–1352.

Holmstrom, B. and Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *the Quarterly Journal of economics*, 112(3):663–691.

Humlum, A. (2022). Robot adoption and labor market dynamics. Working Paper.

Hummels, D., Jørgensen, R., Munch, J., and Xiang, C. (2014). The wage effects of offshoring: Evidence from danish matched worker-firm data. *American Economic Review*, 104(6):1597–1629.

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43.

Jeenas, P. (2024). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. Working Paper.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Kiyotaki, N. and Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2):211–248.

Koch, M., Manuylov, I., and Smolka, M. (2021). Robots and firms. *The Economic Journal*, 131(638):2553–2584.

Kroen, T., Liu, E., Mian, A. R., and Sufi, A. (2024). Falling rates and rising superstars. National Bureau of Economic Research.

Leigh, N. G. and Kraft, B. R. (2018). Emerging robotic regions in the united states: insights for regional economic evolution. *Regional Studies*, 52(6):804–815.

Lian, C. and Ma, Y. (2021). Anatomy of corporate borrowing constraints. *The Quarterly Journal of Economics*, 136(1):229–291.

Linn, M. and Weagley, D. (2024). Uncovering financial constraints. *Journal of Financial and Quantitative Analysis*, 59:2582–2617.

Liu, E., Mian, A. R., and Sufi, A. (2023). Low interest rates, market power, and productivity growth. *Econometrica*, 91(3):1015–1053.

Miranda-Agrippino, S. and Ricco, G. (2021). The transmission of monetary policy shocks. *American*

Economic Journal: Macroeconomics, 13(3):74–107.

Ottonello, P. and Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6):2473–2502.

Peters, R. H. and Taylor, L. A. (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272.

Rampini, A. and Viswanathan, S. (2010). Collateral, risk management, and the distribution of debt capacity. *Journal of Finance*, 65:2293–2322.

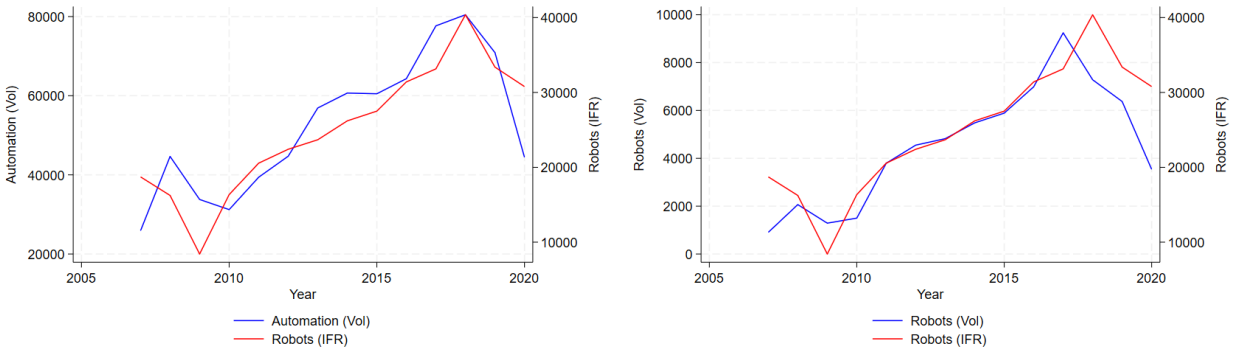
Rampini, A. and Viswanathan, S. (2013). Collateral and capital structure. *Journal of Financial Economics*, 109:466–492.

Rodrigo, R. (2023). Robot adoption, organizational capital and the productivity paradox. *Working Paper*.

Swanson, E. and Jayawickrema, V. (2024). Speeches by the fed chair are more important than fomc announcements: An improved high-frequency measure of u.s. monetary policy shocks. Working Paper.

Wang, J. (2022). Essays on trade, technology, and banking. Working Paper.

Whited, T. (1992). Debt, liquidity constraints, and corporate investment: Evidence from panel data. *Journal of Finance*, 47:1425–1460.



(a) Automation vs. IFR

(b) Robots vs. IFR

Figure 3: Comparison of Imports Data vs. IFR Data

These figures show, in blue, all transactions of automation technologies *(Panel A)* and industrial robots *(Panel B)* imported in volume in 2007-2019, excluding logistic firms, and compare them with the total stock of industrial robots according to the IFR, in red.

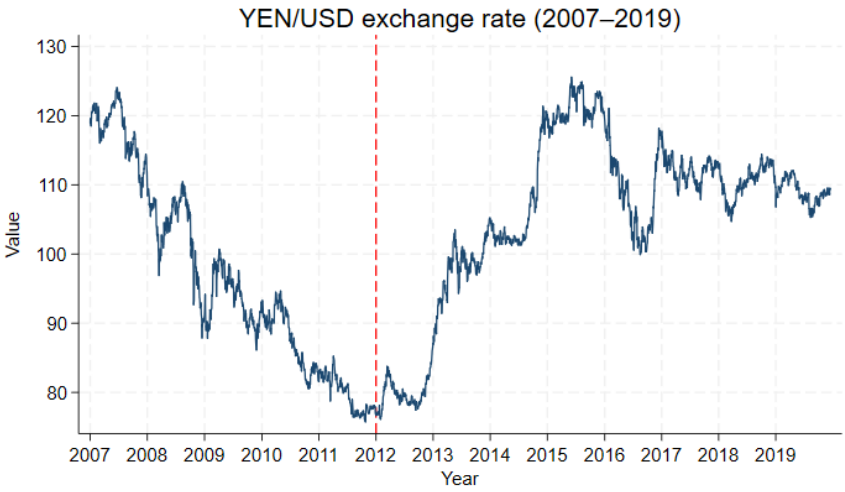


Figure 4: Japanese Market Share and the Yen-USD Exchange Rate

These figures show the Yen-USD exchange rate (per 1 USD), along with the market share of Japanese imports of automation to the United States. The market share is defined as the share of total automation related imports that are sourced from Japan.

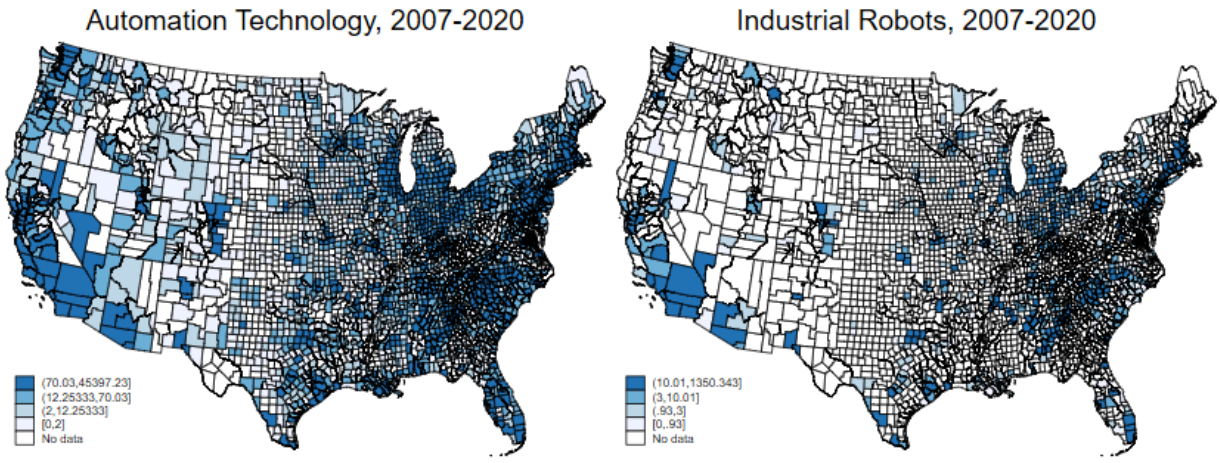
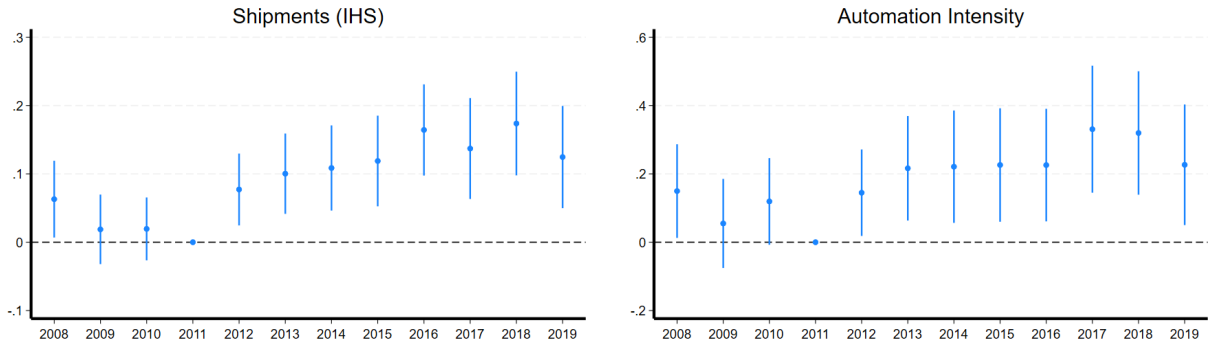


Figure 5: Geographic Distribution of Automation and Robots, 2007 - 2020
These figures shows the maps with all transactions of automation technologies *(Panel A)* and industrial robots imported *(Panel B)* in volume in 2007-2019 by county, excluding imports from logistic firms.

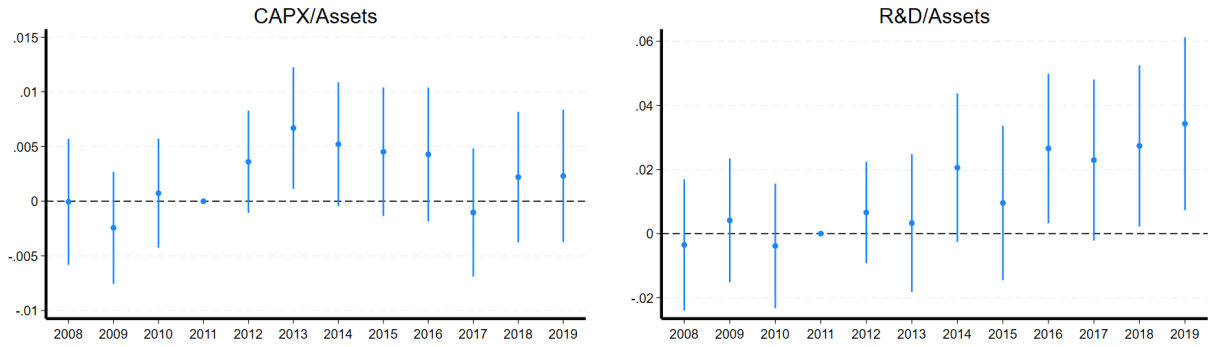


(a) Number of shipments

(b) Automation intensity

Figure 6: Differential Automation Response to YEN-USD Depreciation: High vs. Low Cash Constraints

This figure plots the dynamic response of the inverse hyperbolic sine of automation, in Panel (a) and automation intensity, in Panel (b), to a reduction in robot prices according to equation 5.3. The sample considers all manufacturing firms from 2007 to 2019. Firm exposure to robot prices is measured as the year relative to 2011 (before the large YEN-USD depreciation) interacted with an indicator of firms' high cash constraints, based on each firm's cash liquidity in the first year it appears in the data relative to the median firm for that year. All regressions include firm controls, as well as firm-fixed effects. The lines indicate the 95% confidence intervals using standard errors clustered at the firm level.

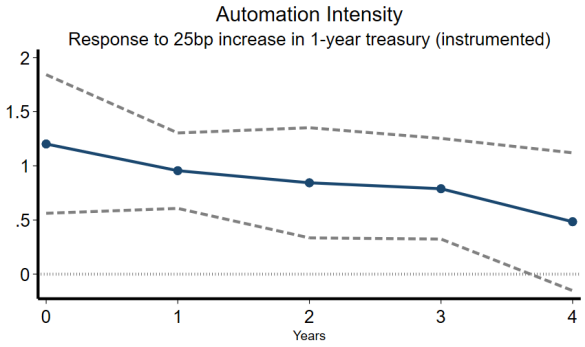


(a) Capital expenditure by assets

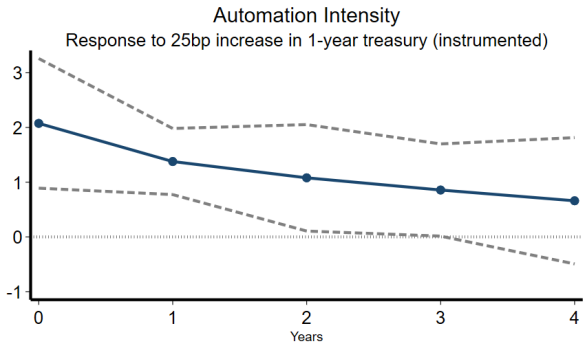
(b) R&D by assets

Figure 7: Differential Investment Response to YEN-USD Depreciation: CAPX and R&D by Cash Constraints

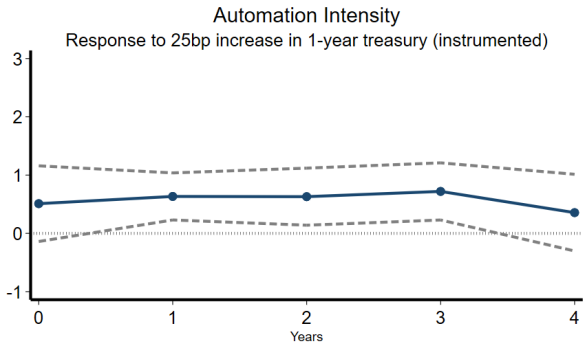
This figure plots the dynamic response of capital expenditure (CAPX) scaled by assets in Panel (a), and of R&D spending scaled by assets in Panel (b), to the Japanese Yen depreciation. The sample includes all manufacturing firms from 2007 to 2019. We measure firm exposure to the exchange rate shock as the year relative to 2011 (before the large yen depreciation began), interacted with an indicator of high cash constraints based on each firms' cash liquidity levels in the first year they appear in the data relative to the median firm for that year. All regressions include firm controls, as well as firm-fixed effects. The lines indicate the 95% confidence intervals using standard errors clustered at the firm level.



(a) All firms



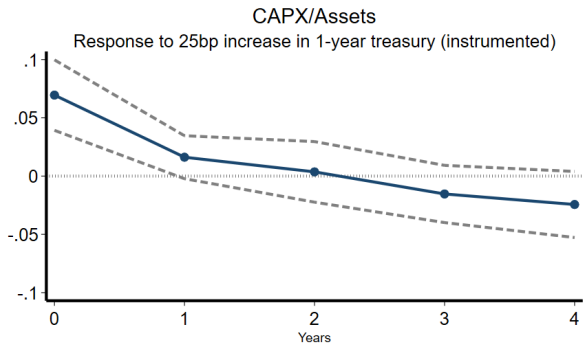
(b) Firms with high cash constraints



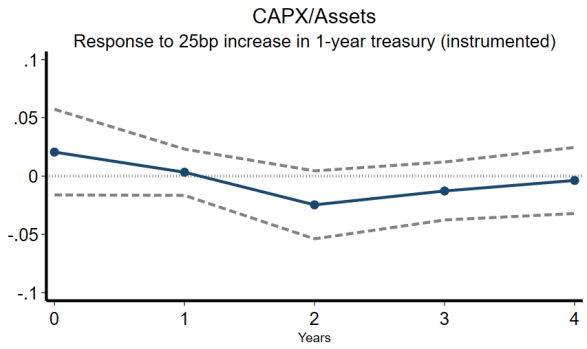
(c) Firms with low cash constraints

Figure 8: Differential Automation Response to Monetary Policy Shocks: By Cash Constraints

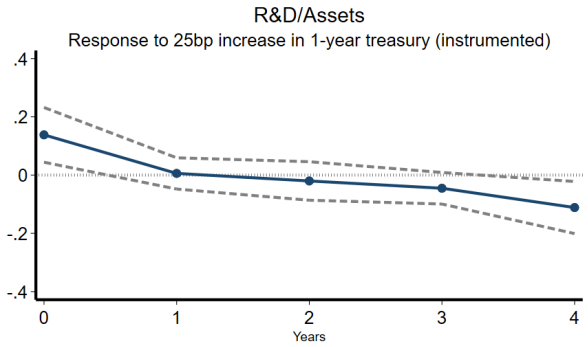
This figure plots the dynamic response of investment in automation technology (scaled by assets) to a decrease in the 1-year Treasury rate estimated using equation 5.4, instrumented by cumulative high-frequency monetary policy shocks. The sample includes all manufacturing firms from 2007 to 2019. Panel (a) considers all firms. Panel (b) and (c) split the sample into high- and low-cash constraints, based on firms' cash liquidity levels in the first year they appear in the data relative to the median firm for that year. Each point represents the estimated coefficient instrumented of the 1-year Treasury rate. All regressions include firm and macro (lagged inflation, GDP growth, and industrial production) controls, as well as firm-fixed effects. The dashed lines indicate 95% confidence intervals with standard errors clustered at the firm level.



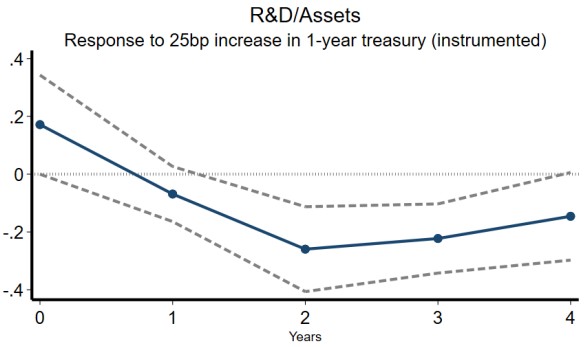
(a) Firms with high cash constraints



(b) Firms with low cash constraints



(c) Firms with high cash constraints



(d) Firms with low cash constraints

Figure 9: Differential Investment Response to Monetary Policy Shocks, CAPX and R&D By Cash Constraints

Panels (a) and (b) plot the dynamic response of capital expenditure (scaled by assets) and Panels (c) and (d) show the response of R&D expenditure (scaled by assets) to a decrease in the 1-year Treasury rate estimated using equation 5.4, instrumented by cumulative high-frequency monetary policy shocks, splitting the sample into high- and low-cash constraints, based on firms' cash liquidity levels in the first year they appear in the data relative to the median firm for that year. Each point represents the estimated coefficient instrumented of the 1-year Treasury rate. All regressions include firm and macro (lagged inflation, GDP growth, and industrial production) controls, as well as firm-fixed effects. The dashed lines indicate the 95% confidence intervals using standard errors clustered at the firm level.

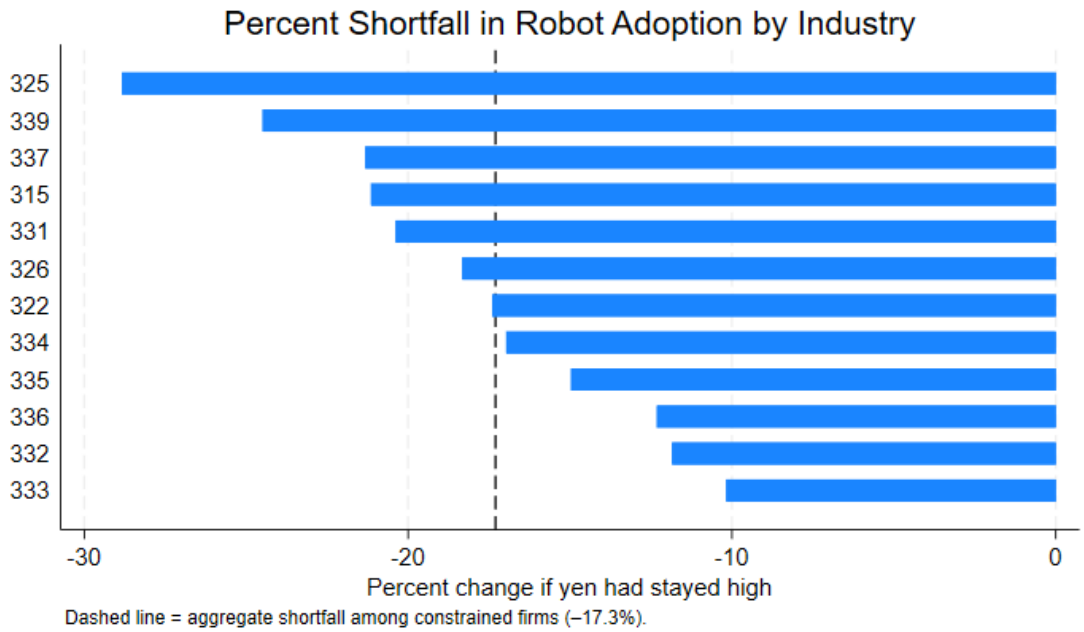


Figure 10: Counterfactual: Automation Without the YEN Depreciation, by 3-digit Industry

The figure plots by industry how many robots financially constrained firms would have adopted if the YEN-USD exchange rate would have remained at its 2011 level. We hold the yen constant at the 2011 average to eliminate the robot price declines associated with the depreciation. Using the estimated coefficient on the $HighCashCon \times YEN - USD$ interaction, we predict robot imports for 2012 to 2015 for 3-digit NAICS industries with more than 20 imported robots.

Table 3: Top Suppliers of Automation Technologies and Industrial Robots

This table examines the geographic and firm-level concentration of automation technology suppliers to the United States. Panel A presents the top five countries by share of total automation technology exports to the U.S., while Panel B focuses specifically on the top five countries supplying industrial robots. Both panels report each country's share of total import volume and total shipments. Panel C identifies the top ten individual companies supplying automation technologies and industrial robots to U.S. firms, along with their countries of origin. The sample includes all import transactions in our dataset, excluding only those conducted by logistics intermediaries to focus on end-user adoption by operating companies.

Panel A: Automation by Country of Origin

Country	Volume (%)	Shipments (%)
Japan	21.2	20.28
China	18.99	23.01
Germany	14.54	10.43
Taiwan	8.17	9.8
South Korea	8.05	6.05
Others	29.05	30.43
Total	*100*	*100*

Panel B: Robots by Country of Origin

Country	Volume (%)	Shipments (%)
Japan	57.66	39.76
Germany	9.24	9.61
Austria	6.50	6.91
China	5.12	12.81
France	4.87	9
South Korea	4.14	4.45
Others	12.47	17.46
Total	*100*	*100*

Panel C: Top Suppliers

Automation Technologies		Industrial Robots	
Supplier	Country	Supplier	Country
Fanuc Corporation	Japan	Fanuc Corporation	Japan
Daimler AG	Germany	Yaskawa Electric Corp.	Japan
Toyota Motor Corporation	Japan	Engel Austria GmbH	Austria
Samsung Electronics Co. Ltd.	S. Korea	Sepro Robotique	France
Doosan Infracore Co. Ltd.	S. Korea	Kawasaki Heavy Industries	Japan
Hyundai Wia Corporation	S. Korea	Kuka Roboter GmbH	Germany
Aisin AW Co. Ltd.	Japan	Nachi Fujikoshi Corp.	Japan
Yamazaki Mazak	Japan	ABB Engineering (SH) Ltd.	China
Doosan Machine Tools	S. Korea	Yushin Precision Equip	Japan
Volkswagen AG	Germany	Comau S. p.A.	Italy

Table 4: Who adopts Automation?

This table presents the distribution of automation technology and industrial robot adoption across industry sectors for U.S. publicly traded firms matched to Compustat from 2007-2019. Panel A shows adoption patterns across broad sectors, reporting both within-sector adoption rates (the percentage of firms within each sector that adopt automation) and between-sector shares (the percentage of total adopters that come from each sector), including a more detailed breakdown for the manufacturing subsector. The sample excludes firms in finance and utilities sectors and those identified as logistics intermediaries.

Panel A: Industry Breakdown

	Automation Technologies		Industrial Robots	
Sector	Within	Between	Within	Between
Agriculture, forestry, fishing and mining	0.09	5.05	0.01	3.06
Wholesale and retail	0.27	13.31	0.04	12.71
Services	0.07	10.43	0.01	6.19
Manufacturing	0.24	71.21	0.07	78.04
Metal products and machinery	0.48	17.44	0.08	20.99
Computers and electronics	0.25	20.62	0.04	20.8
Automotive	0.48	8.54	0.15	17.41
Other manufacturing	0.15	24.61	0.02	18.84
Total		*100*		*100*

Panel B: Top Importers of Automation

Company	Country of Origin	Description
Beckman Coulter, Inc.	Japan	Manufacture of medical equipment
Thermo Fisher Scientific Inc.	Germany	Manufacture of clinical instruments
Intel Corporation	Japan	Design and manufacture computer components
General Motors Company	Japan	Automobile manufacturer
Lear Corporation	Germany	Manufacture of car seats and electrical systems
Kimberly-Clark Corporation	Hong Kong	Consumer goods and personal care
Hamilton Beach Brands, Inc.	China	Home appliances and restaurant equipment
True Value Company	China	Manufacture of Home products, tools and hardware
Hardinge Inc.	Taiwan	Machine tool builder
Mettler-Toledo Int. Inc.	China	Manufactures and sells precision instruments
Caterpillar Inc.	China	Construction equipment

Panel C: Top Importers of Industrial Robots

Company	Country of Origin	Description
iRobot Corporation	Hong Kong	Design and builds consumer robots
Amazon.com, Inc.	China	e-commerce
Abbott Laboratories	Switzerland	Medical devices and health care
Triangle Brick Company	Germany	Brick manufacturer
Hanesbrands Inc.	Vietnam	Clothing company
Ford Motor Company	Germany	Automobile manufacturer
Axcelis Technologies, Inc.	Japan	Components for semiconductors.
Berry Plastics Corporation	Germany	Plastic Packaging Manufacturer
Sears Holdings Corporation	Hong Kong	Department stores
Beckman Coulter, Inc.	Hong Kong	Healthcare diagnosis

Table 5: Automation Adoption Across Firms: Baseline Characteristics

This table examines the role of firm characteristics in their automation adoption decisions. The sample contains publicly listed firms from Compustat, excluding those in the finance and utility sectors. Panel A shows the results for our baseline specification described in equation 5.1, for the following measures of automation adoption as dependent variables: Col. 1 uses an indicator variable for the adoption of automation technologies, *Automate*; Col. 2 uses the the number of shipments, *Shipments* (IHS); Col. 3 uses a ratio of the number of shipments to assets, *AutomationIntensity*; Col. 4 uses the stock of shipments upto time t, $Shipments - Stock$ (IHS); Col. 5 uses the volume of imported automation in Twenty-Foot Equivalent Units (TEU), $Volume(IHS)$. We use the inverse hyperbolic sine transformation of Shipments, Shipments-stock, and Volume variables to acccommodate the skewed distribution. Panel B presents results for the Manufacturing Sector (NAICS 31-33). All regressions are estimated using ordinary least squares with industry x year fixed effects and standard errors clustered at the firm level are reported in parentheses. The Appendix contains detailed definitions of all the variables. ****$p < 0.01$, **$p < 0.05$, *$p < 0.1$

Panel A: Baseline Characteristics

	(1)	(2)	(3)	(4)	(5)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Shipments-Stock (IHS)	Volume (IHS)
Log(Assets)_{t-1}	0.035***	0.050***	0.033***	0.160***	0.055***
	(0.001)	(0.002)	(0.004)	(0.007)	(0.003)
Log(Age)_{t-1}	0.030***	0.044***	0.049***	0.167***	0.044***
	(0.003)	(0.004)	(0.009)	(0.012)	(0.004)
Leverage_{t-1}	0.009***	0.013***	0.000	0.038***	0.016***
	(0.001)	(0.002)	(0.004)	(0.007)	(0.002)
ROA_{t-1}	-0.009***	-0.014***	-0.004**	-0.044***	-0.015***
	(0.001)	(0.001)	(0.002)	(0.004)	(0.001)
FE	Ind x year	Ind x year	Ind x year	Ind x year	Ind x year
N	40,186	39,694	39,765	40,186	39,693
Adj. R-sq	0.209	0.169	0.070	0.355	0.144

Panel B: Manufacturing Sector

	(1)	(2)	(3)	(4)	(5)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Shipments-Stock (IHS)	Volume (IHS)
Log(Assets)_{t-1}	0.052***	0.076***	0.056***	0.232***	0.084***
	(0.002)	(0.004)	(0.007)	(0.011)	(0.004)
Log(Age)_{t-1}	0.034***	0.052***	0.064***	0.191***	0.054***
	(0.004)	(0.006)	(0.014)	(0.017)	(0.007)
Leverage_{t-1}	0.013***	0.020***	0.004	0.057***	0.023***
	(0.002)	(0.003)	(0.005)	(0.010)	(0.003)
ROA_{t-1}	-0.013***	-0.019***	-0.007***	-0.059***	-0.022***
	(0.001)	(0.002)	(0.003)	(0.006)	(0.002)
FE	Ind x year	Ind x year	Ind x year	Ind x year	Ind x year
N	21,974	21,568	21,635	21,974	21,586
Adj. R-sq	0.251	0.203	0.076	0.412	0.174

Table 6: Automation and the Exchange Rate Shock

This table examines the effects of Japanese Yen depreciation on firm-level automation adoption. The dependent variable is one of the following: *Automate* is a indicator varible equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. *Yen-USD* measures the number of Japanese Yen required to purchase one U.S. dollar; an increase indicates Yen depreciation, making Japanese robots cheaper for U.S. importers. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Yen Depreciation and Automation Adoption

	(1)	(2)	(3)
	Automate Dummy	Shipments (IHS)	Automation Intensity
$YEN - USD_t$	0.225***	0.424***	0.525***
	(0.044)	(0.074)	(0.178)
FE	Firm	Firm	Firm
Firm Controls	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes
N	21,637	21,230	21,295
adj. R-sq	0.489	0.508	0.354

Panel B: Exchange Rate Shocks and Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Yen - USD_t$	0.161***	0.284***	0.524***			
	(0.039)	(0.068)	(0.164)			
$High\ FC_i \times Yen - USD_t$				0.139***	0.201***	0.363**
				(0.042)	(0.073)	(0.174)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	21,637	21,230	21,295	15,049	14,720	14,765
adj. R-sq	0.491	0.511	0.355	0.485	0.520	0.358

Table 7: Automation and the Exchange Rate Shock: Propensity Score Matching
This table examines specifications with propensity scores. estimate each firm's likelihood of being classified as financially constrained via a logit model based on lagged covariates such as size, age, profitability and Tobin's Q, controlling for industry and year effects. This is done on firm-level Compustat data from 2002-2007 – we can use the earlier sample as robot adoption is not used here. We match the average (pre-sample) firm propensity score with the same firm in the baseline sample, matching about three-quarters of firms (some firms are not in both samples). The weight for a constrained (treated) observation is $(1/\hat{p})$; for an unconstrained observation it is $1/(1-\hat{p})$. To limit variance inflation, we (i) inspect the weight distribution and (ii) verify post-weighting balance with standardized mean differences. Once we have created propensity scores for firms' constraints, we conduct two tests. First, in Panel A, we weight the baseline specification using the inverse propensity weights. This method preserves all observations but changes each observation's influence so the treated and control groups are comparable on observables in expectation. Second, we match each constrained firm to its nearest unconstrained counterpart with a similar estimated propensity score. This matched sample restricts attention to firms that are observationally similar in size, profitability, and Tobin's Q prior to the shock. We restrict the analysis to pairs that lie on common support, ensuring direct like-with-like comparisons and run the same fixed-effects regression on the matched sample in Panel B. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Inverse Prop. Weighting (IPW)

	Cash Constraints			Financial Constraints		
	Automation Dummy	Shipments (IHS)	Shipments/ Assets	Automation Dummy	Shipments (IHS)	Shipments/ Assets
$High\ CashCon_i \times YEN-USD_t$	0.169***	0.286***	0.460**			
	(0.054)	(0.099)	(0.232)			
$High\ FC_i \times YEN-USD_t$				0.129**	0.188**	0.352
				(0.052)	(0.091)	(0.217)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
N	15,138	14,767	14,847	14,174	13,853	13,902
adj. R-sq	0.486	0.516	0.338	0.479	0.523	0.339

Panel B: Matched Sample (on Constraints)

	Cash Constraints			Financial Constraints		
	Automation Dummy	Shipments (IHS)	Shipments/ Assets	Automation Dummy	Shipments (IHS)	Shipments/ Assets
$High\ CashCon_i \times YEN-USD_t$	0.141**	0.219*	0.588**			
	(0.063)	(0.114)	(0.269)			
$High\ FC_i \times YEN-USD_t$				0.116*	0.104	0.302
				(0.062)	(0.113)	(0.266)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
N	8,351	8,129	8,162	7,869	7,667	7,692
adj. R-sq	0.485	0.511	0.362	0.484	0.516	0.365

Table 8: Automation and Monetary Policy Shock

This table explores the effects on automation of changes in the 1-year Treasury rate, instrumented by Swanson and Jayawickrema (2024) orthogonalized monetary policy surprises. The dependent variable is one of the following: *Automate* is a indicator variable equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions control for firm and macro controls, and include firm and year fixed effects. Standard errors clustered at the firm level are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Monetary Policy Shocks

	(1)	(2)	(3)
	Automate Dummy	Shipments (IHS)	Automation Intensity
$Tyield_{t-1}$	-0.090***	-0.179***	-0.192***
	(0.019)	(0.030)	(0.072)
Firm Controls	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes
FE	Firm	Firm	Firm
N	19,600	19,221	19,297
First-stage F-stat	2,242.155	2,170.909	2,096.904

Panel B: Monetary Policy Shocks and Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Tyield_{t-1}$	-0.085***	-0.150***	-0.268***			
	(0.023)	(0.038)	(0.095)			
$High\ FC_i \times Tyield_{t-1}$				-0.074***	-0.119***	-0.122
				(0.023)	(0.037)	(0.094)
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
N	19,600	19,221	19,297	13,235	12,930	12,989
First-stage F-stat	6,152.224	5,847.410	5,781.366	5,532.318	5,362.188	5,281.104

Table 9: Automation Response to Combined Exchange Rate and Monetary Policy Shocks

This table examines the joint effects of exchange rate and monetary policy shocks on firm-level automation investment. The Yen-USD exchange rate (divided by 100 and logged) captures exogenous variation in robot prices, while the 1-year Treasury rate is instrumented using Swanson and Jayawickrema (2024) orthogonalized monetary policy surprises to identify financing cost shocks. The dependent variables are: *Automate*, an indicator equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, the inverse hyperbolic sine transformation of annual automation shipments; and $AutomationIntensity$, automation shipments scaled by lagged total assets. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions include firm fixed effects with standard errors clustered at the firm level. Columns with year fixed effects control for aggregate trends, while columns without year fixed effects include macroeconomic controls (inflation, industrial production, GDP growth). The sample is restricted to manufacturing firms from 2007-2019.

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times YEN-USD_t$	0.213***	0.369***	0.680***			
	(0.046)	(0.081)	(0.202)			
$High\ CashCon_i \times Tyield_{t-1}$	-0.045***	-0.082***	-0.141**			
	(0.017)	(0.027)	(0.068)			
$High\ FC_i \times YEN-USD_t$				0.178***	0.251***	0.428**
				(0.050)	(0.088)	(0.217)
$High\ FC_i \times Tyield_{t-1}$				-0.042**	-0.074***	-0.044
				(0.017)	(0.027)	(0.067)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	19,600	19,221	19,297	13,235	12,930	12,989
First-stage F-stat	19,044.840	17,763.284	17,652.056	29,235.245	27,088.671	26,922.842

Table 10: Robustness: Alternate Measures of Constraints

This table examines the joint effects of exchange rate and monetary policy shocks on firm-level automation investment. The Yen-USD exchange rate (divided by 100 and logged) captures exogenous variation in robot prices, while the 1-year Treasury rate is instrumented using Swanson and Jayawickrema (2024) orthogonalized monetary policy surprises to identify financing cost shocks. The dependent variables are: $Automate$, an indicator equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, the inverse hyperbolic sine transformation of annual automation shipments; and $AutomationIntensity$, automation shipments scaled by lagged total assets. Panel A examines heterogeneity by debt constraints, where $High\ FC$ (HM) indicates firms above the median Hoberg and Maksimovic (2015) financing constraint measure in their first sample year. These constraints are only available until 2015, but firms maintain their initial classification through 2019. Panel B examines heterogeneity by Leverage ratios defined as Short-term & Long-term Debt over Total Assets, with firms above the median classified as ($High\ Leverage$). All regressions include firm fixed effects with standard errors clustered at the firm level. Columns with year fixed effects control for aggregate trends, while columns without year fixed effects include macroeconomic controls (inflation, industrial production, GDP growth). The sample is restricted to manufacturing firms from 2007-2019.

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ FC(HM)_i \times YEN - USD_t$	0.150***	0.199*	0.464*			
	(0.058)	(0.102)	(0.247)			
$High\ FC(HM)_i \times Tyield_{t-1}$	-0.036*	-0.067**	-0.135*			
	(0.021)	(0.033)	(0.079)			
$High\ Leverage_i \times YEN - USD_t$				0.175***	0.337***	0.546***
				(0.047)	(0.083)	(0.205)
$High\ Leverage_i \times Tyield_{t-1}$				-0.054***	-0.099***	-0.117*
				(0.017)	(0.027)	(0.067)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	11,584	11,322	11,349	18,341	17,966	18,052
First-stage F-stat	19,044.840	17,763.284	17,652.056	29,235.245	27,088.671	26,922.842

Table 11: Robustness: Controlling for Outsourcing

This table examines the joint effects of exchange rate and monetary policy shocks on firm-level automation investment. The Yen-USD exchange rate (divided by 100 and logged) captures exogenous variation in robot prices, while the 1-year Treasury rate is instrumented using Swanson and Jayawickrema (2024) orthogonalized monetary policy surprises to identify financing cost shocks. The dependent variables are: $Automate$, an indicator equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, the inverse hyperbolic sine transformation of annual automation shipments; and $AutomationIntensity$, automation shipments scaled by lagged total assets. $Outsourcing/dummy$ is an indicator variable equal to 1 if the firm outsources (defined as importing products (not including automation) in its industry) and $OutsouringIntensity$ is the number of outsourced shipments scaled by lagged value of assets. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions include firm and year fixed effects with standard errors clustered at the firm level. The sample is restricted to manufacturing firms from 2007-2019.

Panel A: Response to both Shocks by Cash Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Yen_t$	0.207***	0.360***	0.665***	0.200***	0.350***	0.617***
	(0.045)	(0.081)	(0.201)	(0.046)	(0.081)	(0.200)
$High\ CashCon_i \times Tyield_t$	-0.042**	-0.078***	-0.132*	-0.042**	-0.082***	-0.124*
	(0.017)	(0.027)	(0.068)	(0.017)	(0.027)	(0.067)
Outsourcing dummy	0.086***	0.110***	0.199***			
	(0.009)	(0.013)	(0.033)			
Outsourcing intensity				0.000***	0.000***	0.001***
				(0.000)	(0.000)	(0.000)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	19,600	19,221	19,297	19,258	18,882	18,989

Panel B: Response to both Shocks by Debt Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS	Automation Intensity
$HighFC_i \times Yen_t$	0.173***	0.245***	0.418*	0.167***	0.237***	0.383*
	(0.050)	(0.088)	(0.216)	(0.050)	(0.088)	(0.213)
$High\ FC_i \times Tyield_{t-1}$	-0.040**	-0.071***	-0.039	-0.041**	-0.074***	-0.035
	(0.017)	(0.027)	(0.067)	(0.017)	(0.027)	(0.066)
Outsourcing dummy	0.078***	0.097***	0.185***			
	(0.010)	(0.015)	(0.037)			
Outsourcing intensity				0.000***	0.000***	0.001***
				(0.000)	(0.000)	(0.000)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	13,235	12,930	12,989	12,964	12,661	12,741

Table 12: Robustness: Star Firms and Investment Opportunities

In Panel A, we present robustness tests by dropping all "star" firms, defined as top 5% of firms by Market Capitalization within an industry(NAICS3)-year following **?**. In Panel B, we repeat the specification in Table 9 with the following additional firm-level controls: Tobin's Q is an estimate of Tobin's Q adjusted for intangible capital from Peters and Taylor (2017); Intangible Intensity defined as the value of intangible capital (from Peters and Taylor (2017)) over total capital; and Sales growth defined as annual log difference in firm sales. The dependent variables are: *Automate*, an indicator equal to 1 if the firm imports automation equipment in that year; and *AutomationIntensity*, automation shipments scaled by lagged total assets. *High CashCon* and *HighFC* are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Drop Star Firms

	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times YEN - USD_t$	0.213***	0.376***	0.688***			
	(0.047)	(0.084)	(0.209)			
$High\ CashCon_i \times Tyield_{t-1}$	-0.041**	-0.079***	-0.129*			
	(0.018)	(0.028)	(0.071)			
$High\ FC_i \times YEN - USD_t$				0.174***	0.247***	0.418*
				(0.050)	(0.089)	(0.220)
$High\ FC_i \times Tyield_{t-1}$				-0.040**	-0.072***	-0.039
				(0.017)	(0.027)	(0.068)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	18,981	18,602	18,682	13,104	12,799	12,858

Panel B: Additional Controls

	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Yen - USD_t$	0.190***	0.315***	0.704***			
	(0.059)	(0.107)	(0.273)			
$High\ CashCon_i \times Tyieldt - 1$	-0.047**	-0.081**	-0.156*			
	(0.022)	(0.035)	(0.091)			
$High\ FC_i \times Yen - USD_t$				0.160**	0.194*	0.350
				(0.063)	(0.113)	(0.285)
$High\ FC_i \times Tyieldt - 1$				-0.042*	-0.072**	-0.008
				(0.022)	(0.034)	(0.087)
$Tobins\ Q_{t-1}$	-0.007***	-0.009**	-0.002	-0.009**	-0.011*	-0.001
	(0.003)	(0.004)	(0.011)	(0.004)	(0.006)	(0.016)
$Intangible\ Intensity_{t-1}$	0.126*	0.166	-0.154	0.168*	0.227	-0.359
	(0.075)	(0.141)	(0.352)	(0.087)	(0.161)	(0.379)
$Sales\ Growth_{t-1}$	0.013***	0.019***	0.027*	0.015***	0.022***	0.030*
	(0.004)	(0.006)	(0.014)	(0.005)	(0.007)	(0.018)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	12,605	12,231	12,350	10,121	9,820	9,910
First-stage F-stat	17,310.793	15,982.222	15,856.311	18,148.538	16,453.727	16,741.497

INTERNET APPENDIX

Automation Under Constraints: Exchange Rates, Interest Rates, and Firm Investment

Meghana Ayyagari Vojislav Maksimovic Rodimiro Rodrigo Ariel Weinberger

This Internet Appendix provides supplemental materials referenced in the main text. Section IA.1 defines all variables used in the analysis, including their precise construction and data sources. Specifically it describes the procedures used to merge and clean the transaction-level Panjiva import data with Compustat firm-year financial data; summary statistics of key variables; and definitions of automation technologies used in the paper. Section IA.2 reports additional robustness tests and alternative specifications discussed in Section 5 of the paper. Section IA.3 presents proofs to the main propositions as well as theoretical extensions.

IA.1 Variable Definitions, Sample Construction, and Summary Statistics

Table IA.1.1: Variable Definitions and Sources

Variable	Definition and Construction	Data Source
Automate	Indicator = 1 if firm i imports at least one automation-related product (HS codes in Table IA.3) in year t.	Panjiva
Shipments (IHS)	Inverse hyperbolic sine (IHS) of the annual number of automation-related import shipments per firm-year.	Panjiva
Automation Intensity	Number of automation shipments divided by lagged total assets.	Panjiva + Compustat
Shipments–Stock (IHS)	IHS of the cumulative number of automation shipments up to year t.	Panjiva
Volume (IHS)	IHS of total container volume (TEU) of automation imports per firm-year.	Panjiva
Log(Assets)	Natural logarithm of total assets (Compustat item AT).	Compustat
Log(Age)	Log of years since first appearance in Compustat.	Compustat
Leverage	Total debt (Compustat item DLC + DLTT) / Total assets (AT).	Compustat
ROA	Net income (Compustat item NI) / Total assets (AT).	Compustat
Cash/Assets	Cash and short-term investments (Compustat item CHE) / Total assets (AT)	Compustat
High Cash Constr	= 1 if Cash/Assets below median in firm's first sample year (internal liquidity constraint indicator).	Derived
High Leverage	= 1 if Leverage above median in firm's first sample year (debt constraint indicator).	Derived
High FC (LW)	= 1 if firm above median in Linn & Weagley (2024) predicted financial constraints index.	Linn and Weagley (2024)
High FC (HM)	= 1 if firm above median in Hoberg–Maksimovic (2015) text-based debt constraint measure.	Hoberg and Maksimovic (2015)
YEN–USD	Log yen–per–USD exchange rate (divided by 100); A higher value indicates a yen depreciation, corresponding to lower U.S. dollar prices of Japanese robots and automation equipment.	FRED (Federal Reserve Economic Data)
BroadUSD	Log trade-weighted U.S. Dollar Index (base 100).	FRED
RMB–USD	Log of the Chinese yuan–per–U.S.-dollar nominal exchange rate (divided by 100). An increase indicates a yuan depreciation, serving as a control for price shocks from China-sourced automation or other imported capital goods.	FRED
EME–USD	Log of the Nominal Emerging Market Economies U.S. Dollar Index (divided by 100) from FRED. A higher value represents a stronger U.S. dollar relative to a basket of emerging-market currencies, capturing general movements in dollar strength unrelated to Japan-specific shocks.	FRED
Tyield	1-year U.S. Treasury yield; instrumented by cumulative FOMC monetary policy surprises (Bauer & Swanson 2023).	FRED / FRBSF MPS
Baa–10y Spread	Moody's Baa corporate bond yield minus 10-year Treasury yield — investment-grade credit spread capturing risk premia over risk-free rate (Gilchrist & Zakrajšek 2012).	FRED
Excess Bond Premium (EBP)	Non-default component of corporate bond spreads — residual after removing expected default risk and bond characteristics; proxy for credit-supply or risk-appetite.	Gilchrist & Zakrajšek (2012) via FRED
Dist to Default	Distance-to-default measure based on market leverage and asset volatility.	Compustat / CRSP derived
Outsourcing Dummy	= 1 if firm imports non-automation goods within its own NAICS industry code (outsourcing indicator).	Panjiva
Outsourcing Intensity	Total number of non-automation imports in the firm's NAICS industry scaled by lagged firm assets.	Panjiva
Tobin's Q	Market value of assets / replacement cost, adjusted for intangibles (Peters & Taylor 2017).	Compustat
Intangible Intensity	Intangible capital / total capital (from Peters & Taylor 2017 series).	Compustat
Sales Growth	Annual log change in sales (REVt – REVt-1).	Compustat

Table IA.1.2: Summary Statistics of Key Variables

Variable	N	sd	mean	p25	p50	p75
Automate	21,974	0.359	0.152	0	0	0
Shipments (IHS)	21,568	0.587	0.208	0	0	0
Automation Intensity	21,635	1.299	0.300	0	0	0
Shipments Stock (IHS)	21,974	1.349	0.694	0	0	0.882
Volume (IHS)	21,586	0.686	0.214	0	0	0
YEN-USD	28,533	0.133	0.010	-0.065	0.056	0.099
Tyield	28,533	1.202	1.113	0.223	0.524	1.492
Cumulative MPS	28,533	0.110	-0.230	-0.281	-0.206	-0.182
Log (Assets)	21,974	2.713	5.133	3.369	5.186	7.050
Log (Age)	21,974	1.095	2.534	1.792	2.773	3.332
Leverage	21,974	1.009	0.347	0.001	0.154	0.343
ROA	21,974	2.333	-0.552	-0.340	-0.002	0.064
High CashCon	28,533	0.486	0.384	0	0	1
High FC	17,673	0.500	0.500	0	1	1
High Leverage	25,770	0.498	0.456	0	0	1
High FC (HM)	15,452	0.499	0.467	0	0	1

Table IA.1.3: Data Processing and Sample Construction

Step	Description	Observations
1	Panjiva trade transactions (one shipment): keep only HS codes with industrial robots or automation HS6 codes as defined in Acemoglu and Restrepo (2022). Keep only if consignee country is listed (and is the United States). We use data from 2007-2019.	490,706 transactions, made up of 61,839 unique firm names (note, many do not have matches outside of Panjiva).
2	Exclude SIC code 4, any firm with name that matches a logistic/distribution/freight firm, as well as robot integrators.	463,609 transactions, 60,737 unique firm names.
3	Use BECRS to match the Panjiva ID to a possible GVKEY	155,740 transactions are matched as having a GVKEY. 8,315 unique Panjiva IDs and 3,551 unique GVKEYs.
4	Collapse transactions to total for each firm-year: now at GVKEY–Year aggregation.	13,353 GVKEY-year observations.
5	Match to Compustat Annual sample. Compustat data downloaded January 2025, using only USA firms with format "INDL". We drop SIC sectors between 40-49, 60-67, and 90-99	5,442 GVKEY-year importers matched to Compustat. Total of 53,000 observations including non-importing Compustat firms.

Table IA.1.4: Automation Technologies, HS 2017

Automation technology	HS 2017 code
Industrial robots	847950
Dedicated machinery	847989
Numerically controlled machines	845710; 845811; 845891; 845921-845961; 846012-846023; 846024-846931; 846221-846241
Automatic machine tools	845611-845690; 846120-846150; 846320; 846330; 846510
Automatic welding machines	851521; 851531
Weaving and knitting machines	844600-844699; 844700-844799
Textile dedicated machinery	844400-844590; 845210; 845230; 845310; 845320
Automatic conveyors	842641; 842649; 842710; 842720; 842870
Regulating and control instruments	902500-903299

Source: Build in HS 2017 codes by the authors, based on the definition of automation technologies in Acemoglu and Restrepo (2022).

IA.2 Robustness Tests

Table IA.2.1: Automation and the Exchange Rate Shock Robustness: Controlling for other exchange rates

This table examines the effects of Japanese Yen depreciation on firm-level automation adoption. The dependent variable is one of the following: *Automate* is a indicator varible equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. *Yen-USD* measures the number of Japanese Yen required to purchase one U.S. dollar; an increase indicates Yen depreciation, making Japanese robots cheaper for U.S. importers. We divide the exchange rates by 100 and take the logarithm. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Automate Dummy	Automate Dummy	Automate Dummy	Shipments (IHS)	Shipments (IHS)	Shipments (IHS)	Automation Intensity	Automation Intensity	Automation Intensity
$High\ CashCon_i \times Yen - USD_t$	0.165***	0.089**	0.068	0.293***	0.174**	0.145*	0.535***	0.401**	0.371*
	(0.039)	(0.043)	(0.047)	(0.068)	(0.069)	(0.075)	(0.168)	(0.181)	(0.197)
$High\ CashCon_i \times RMB - USD_t$	-0.117			-0.248			-0.287		
	(0.106)			(0.185)			(0.491)		
$High\ CashCon_i \times EME - USD_t$		0.159**			0.244**			0.273	
		(0.067)			(0.116)			(0.293)	
$High\ CashCon_i \times BROAD - USD_t$			0.183**			0.275**			0.304
			(0.076)			(0.130)			(0.326)
FE	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	21,637	21,637	21,637	21,230	21,230	21,230	21,295	21,295	21,295
adj. R-sq	0.491	0.491	0.491	0.511	0.511	0.511	0.355	0.355	0.355

Table IA.2.2: Automation and the Exchange Rate Shock: Selection on Never-Adopters
This table conducts robustness by controlling for the propensity to adopt robots based on observable firm measures. We first estimate a firm-level adoption score, defined as the predicted probability of ever adopting automation, using firm characteristics and industry–year effects. We only use 2007-2010 data to construct an average firm-specific score, and then assign each firm a time-invariant score. This score summarizes all observable factors that systematically distinguish adopters from non-adopters. First, in the first 3 columns of Panel A, we re-estimate the specification restricting the sample to firms above the 25th percentile of the adoption-score distribution, that is, to firms that are potential adopters. This trims away firms whose ex-ante probability of adoption is effectively zero, isolating the intensive margin of investment responses among plausible adopters. In the last 3 columns we include industry x year fixed effects using the baseline sample, to once again attempt to compare firms to their closest firms. Second, in Panel B, we include the adoption score directly in the baseline regression - **both in levels and interacted with the exchange-rate shock** - to absorb any component of the response attributable to inherent adoption likelihood rather than financing. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Conditional on Adoption Score and Industry-Year Effects

	Conditional on Adoption Score			Within Industry-Year		
	Automation Dummy	Shipments (IHS)	Automation Intensity	Automation Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Yen - USD_t$	0.161***	0.286***	0.601***	0.138***	0.247***	0.495**
	(0.051)	(0.090)	(0.221)	(0.049)	(0.088)	(0.199)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Ind×Year	Firm, Ind×Year	Firm, Ind×Year
Sample	Score > p25	Score > p25	Score > p25	All	All	All
N	12,896	12,510	12,621	17,225	16,833	16,910
adj. R-sq	0.465	0.499	0.351	0.491	0.515	0.360

Panel B: Control for Propensity to Adopt

	Automation Dummy		Shipments (IHS)		Automation Intensity	
	No conditions	Score > p25	No conditions	Score > p25	No conditions	Score > p25
$High\ CashCon_i \times Yen - USD_t$	0.109***	0.123**	0.137*	0.168*	0.453**	0.542**
	(0.042)	(0.051)	(0.072)	(0.088)	(0.182)	(0.222)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
N	16,765	12,896	16,375	12,510	16,452	12,621
adj. R-sq	0.491	0.466	0.518	0.502	0.359	0.351

Table IA.2.3: Robustness: Alternate Measures of Monetary Policy Shocks

This table examines how monetary conditions affect firm-level automation adoption. We use two measures: (i) the Baa–10y spread (Moody's Baa corporate yield minus the 10-year U.S. Treasury yield), an investment-grade credit spread capturing risk premia over the risk-free rate; and (ii) the Excess Bond Premium (EBP), the non-default component of corporate bond spreads – the cross-sectional residual after removing expected default risk and bond characteristics – interpreted as a credit-supply/risk-appetite gauge (following Gilchrist and Zakrajšek (2012)). The dependent variable is one of the following: *Automate* is a indicator variable equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Panel A: Moody's Baa-10y Treasury spread

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Baa-10y\ spread_{t-1}$	-0.011***	-0.019***	-0.043***			
	(0.004)	(0.006)	(0.016)			
$High\ FC_i \times Baa-10y\ spread_{t-1}$				-0.011***	-0.014**	-0.000
				(0.004)	(0.007)	(0.017)
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
N	21,637	21,230	21,295	15,049	14,720	14,765
adj. R-sq	0.000	0.001	-0.000	0.001	0.002	-0.001

Panel B: Excess Bond Premium

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times Excess\ Bond\ Premium_{t-1}$	-0.006	-0.010*	-0.027*			
	(0.004)	(0.006)	(0.015)			
$High\ FC_i \times Excess\ Bond\ Premium_{t-1}$				-0.008**	-0.010*	0.007
				(0.004)	(0.006)	(0.015)
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
N	21,637	21,230	21,295	15,049	14,720	14,765
adj. R-sq	0.000	0.001	-0.000	0.001	0.002	-0.001

Table IA.2.4: Robustness: Continuous Measures of Constraints

This table examines the effects of Japanese Yen depreciation on firm-level automation adoption. The dependent variable is one of the following: *Automate* is a indicator varible equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. *Yen-USD* measures the number of Japanese Yen required to purchase one U.S. dollar; an increase indicates Yen depreciation, making Japanese robots cheaper for U.S. importers. We divide the exchange rates by 100 and take the logarithm. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$Cash/Assets_{t-1} \times YEN - USD_t$	-0.515***	-0.935***	-1.392***			
	(0.060)	(0.103)	(0.260)			
$Cash/Assets_{t-1} \times Tyield_{t-1}$	0.155***	0.317***	0.615***			
	(0.037)	(0.059)	(0.149)			
$Cash/Assets_{t-1}$	-0.122***	-0.237***	-0.397***			
	(0.027)	(0.042)	(0.107)			
$FC_{t-1} \times YEN - USD_t$				0.113***	0.214***	0.370***
				(0.039)	(0.061)	(0.140)
$FC_{t-1} \times Tyield_{t-1}$				-0.083*	-0.153**	-0.367**
				(0.047)	(0.072)	(0.161)
FC_{t-1}				0.027	0.053	0.149*
				(0.025)	(0.037)	(0.087)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	18,953	18,573	18,654	15,869	15,558	15,597
First-stage F-stat	425.810	422.422	419.290	179.883	172.269	178.167

Table IA.2.5: Robustness: Distance to Default vs. Cash Constraints

This table examines examines a type of "horse race" between Cash Constraints and "Distance to Default" (Gilchrist and Zakrajšek (2012); Ottonello and Winberry (2020)). To construct the distance to default, we follow the Gilchrist and Zakrajšek (2012), using CRSP data to measure equity volatility. The dependent variables are: *Automate*, an indicator variable equal to 1 if the firm imports automation equipment in that year; *Shipments* (IHS), total shipments transformed with the inverse hyperbolic sine; and *Automation Intensity*, shipments as a share of lagged total assets. *YEN-USD* measures the Japanese Yen per U.S. dollar (logged, divided by 100); increases indicate Yen depreciation. *High CashCon* is an indicator for firms above the median in liquidity constraints. $Tyield_{t-1}$ is the lagged Treasury yield. Columns 2, 4, and 6 report instrumental variables estimates. All regressions include firm fixed effects and control variables (size, age, ROA, and leverage). Standard errors clustered at the firm level are reported in parentheses. The sample is restricted to Manufacturing firms. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	Automate Dummy		Shipments (IHS)		Automation Intensity	
	(1)	(2)	(3)	(4)	(5)	(6)
$High\ CashCon_i \times Yen - USD_t$	0.162***	0.229***	0.281***	0.342**	0.674***	0.629*
	(0.057)	(0.081)	(0.106)	(0.151)	(0.257)	(0.364)
$High\ CashCon_i \times Tyield_{t-1}$		-0.055		-0.082		0.004
		(0.035)		(0.057)		(0.135)
$DistDefault_i$	0.000	0.009	-0.001	0.005	0.034*	-0.060
	(0.005)	(0.011)	(0.009)	(0.018)	(0.018)	(0.047)
$DistDefault_i \times Yen - USD_t$	0.002	0.036	-0.030	0.003	-0.069	-0.227
	(0.032)	(0.044)	(0.059)	(0.080)	(0.141)	(0.198)
$DistDefault_i \times Tyield_{t-1}$		-0.056		-0.060		0.356*
		(0.049)		(0.077)		(0.194)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
N	10,245	8,060	9,894	7,739	10,058	7,912
First-stage F-stat		20.380		18.313		19.536

Table IA.2.6: Automation and the Exchange Rate Shock Robustness: Dropping Auto Sector (NAICS 336)

This table examines the effects of Japanese Yen depreciation on firm-level automation adoption. The dependent variable is one of the following: *Automate* is a indicator varible equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, an intensive margin measure of total shipment in one year transformed with the inverse hyperbolic sine; and $AutomationIntensity$, a share of shipments relative to lagged total assets of the firm. *Yen-USD* measures the number of Japanese Yen required to purchase one U.S. dollar; an increase indicates Yen depreciation, making Japanese robots cheaper for U.S. importers. We divide the exchange rates by 100 and take the logarithm. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions are estimated using ordinary least squares with standard errors clustered at the firm level reported in parentheses. The sample is restricted to Manufacturing firms. The Appendix contains detailed definitions of all the variables. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times YEN - USD_t$	0.209***	0.363***	0.631***			
	(0.048)	(0.085)	(0.212)			
$High\ CashCon_i \times Tyield_{t-1}$	-0.041**	-0.073***	-0.118*			
	(0.018)	(0.027)	(0.071)			
$High\ FC_i \times YEN - USD_t$				0.159***	0.210**	0.323
				(0.052)	(0.090)	(0.223)
$High\ FC_i \times Tyield_{t-1}$				-0.037**	-0.062**	-0.027
				(0.017)	(0.026)	(0.069)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	18,445	18,136	18,166	12,398	12,146	12,172
First-stage F-stat	17,350.150	16,338.415	16,066.980	27,613.822	25,998.970	25,494.250

Table IA.2.7: Robustness: Rising vs Declining Interest Rate Regimes

This table examines the joint effects of exchange rate and monetary policy shocks on firm-level automation investment. The Yen-USD exchange rate (divided by 100 and logged) captures exogenous variation in robot prices, while the 1-year Treasury rate is instrumented using Swanson and Jayawickrema (2024) orthogonalized monetary policy surprises to identify financing cost shocks. The dependent variables are: *Automate*, an indicator equal to 1 if the firm imports automation equipment in that year; $Shipments(IHS)$, the inverse hyperbolic sine transformation of annual automation shipments; and $AutomationIntensity$, automation shipments scaled by lagged total assets. $High\ CashCon$ and $HighFC$ are indicator variables equal to one if the firm is above the median in liquidity constraints (inverse of Cash/Total Assets) or financing constraints (Linn and Weagley (2024)) respectively during the first year the firm appears in our sample period. All regressions include firm and year fixed effects with standard errors clustered at the firm level. The sample is restricted to manufacturing firms and is split into two periods: 2007-2013 when interest rates were declining; and 2014-2019 when interest rates were normalized.

Panel A: Cash Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Automate Dummy	Shipments (IHS)	Shipments (IHS)	Automation Intensity	Automation Intensity
	Declining Yields (2007-2013)	Rate Normalization (2014-2019)	Declining Yields (2007-2013)	Rate Normalization (2014-2019)	Declining Yields (2007-2013)	Rate Normalization (2014-2019)
$High\ CashCon_i \times Tyield_{t-1}$	-0.026*	-0.047	-0.054**	-0.227	-0.095*	-0.322
	(0.014)	(0.143)	(0.021)	(0.236)	(0.057)	(0.566)
FE	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year	Firm, year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	8,957	10,344	8,838	10,072	8,818	10,169
First-stage F-stat	2.40e+05	308.290	2.14e+05	250.869	2.16e+05	275.088

Panel B: Debt Constraints

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate Dummy	Automate Dummy	Shipments (IHS)	Shipments (IHS)	Automation Intensity	Automation Intensity
	Declining Yields (2007-2013)	Rate Normalization (2014-2019)	Declining Yields (2007-2013)	Rate Normalization (2014-2019)	Declining Yields (2007-2013)	Rate Normalization (2014-2019)
$High\ FC_i \times Tyield_{t-1}$	-0.030**	0.060	-0.066***	0.112	-0.047	-0.084
	(0.014)	(0.218)	(0.022)	(0.356)	(0.059)	(0.876)
FE	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year	Firm, Year
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
N	7,194	5,943	7,095	5,727	7,073	5,809

IA.3 Theoretical Appendix

This Appendix is organized as follows. Section IA.3.1 introduces notation, the invoice identity, and the key weights used throughout, and Section IA.3.2 states the standing assumptions. Section IA.3.3 records the competitive pricing condition and the b'–FOC. Section IA.3.4 proves a basic "invoice sufficiency" lemma: when finance binds the scale k' is pinned by the cash identity $k'M = c$ so that all local elasticities run through the invoice M, whereas when finance is slack k' is set by the technology FOC and M is no longer sufficient. Section IA.3.5 derives the cap–regime elasticities used in the main text; Sections IA.3.6 and IA.3.7 derive and sign the interior–regime elasticities, proving Proposition 1. Sections IA.3.8, IA.3.9, and IA.3.10 collect auxiliary results on the shadow value λ, region selection, and the slack benchmark used in the "Role of Technology" discussion. Section IA.3.11 proves Proposition 2 (second–order cross–effects). Section IA.3.12 derives the comparative–static arrows underlying Tables 1 and 2. Section IA.3.13 provides the formal version of the diagnostic Proposition 3. Section **??** discusses Asset-Based and Cash-Flow-Based Lending and the Borrowing Constraint.

IA.3.1 Preliminaries, notation, and positivity

In the finance–binding (cash–poor) region, the budget and borrowing constraints imply the *invoice (own–cash) identity*

$$k'\,M(k',\theta) \;=\; c, \qquad M(k',\theta) \;=\; p_k\Big(1+\frac{h}{p_k}-y(R,\pi,\rho)\,m(\theta)\Big)+\frac{\tau}{2}\,k'. \tag{IA.3.1}$$

Borrowing capacity is

$$b' \;\le\; y(R,\pi,\rho)\,m(\theta)\,p_k\,k',$$

where the *advance rate* and its rate sensitivity (from the one–period non–recourse contract) are

$$y(R,\pi,\rho)=\frac{(1-\pi)\rho}{R-\pi}, \qquad \chi_R \equiv \frac{\partial \ln y}{\partial \ln R} = -\frac{R}{R-\pi} < 0. \tag{IA.3.2}$$

From (IA.3.2), $\partial y/\partial R < 0$, $\partial y/\partial \rho > 0$, and $\chi_R = -R/(R-\pi) < 0$. The pledge frontier is increasing and concave: $m(0)=0$, $m'(\theta)>0$, $m''(\theta)<0$, with cap $0 \le \theta \le \bar\theta$ (equivalently $0 \le m(\theta) \le \bar m \equiv m(\bar\theta)$). We impose the positivity condition

$$1+\frac{h}{p_k}-y\,m(\theta) > 0,$$

so that $M>0$ for all $\theta \in [0,\bar\theta]$.

It is convenient to define the *exposure* and *brake* weights:

$$W_{p_k} \equiv \frac{p_k(1-y\,m(\theta))}{M}, \qquad W_f \equiv \frac{p_k\,y\,m(\theta)}{M}, \qquad W_\theta \equiv \frac{p_k\,y\,\theta\,m'(\theta)}{M}, \qquad \kappa \equiv \frac{\frac{\tau}{2}\,k'}{M}. \tag{IA.3.3}$$

Evaluation convention. In cap–regime formulas we evaluate $m(\cdot)$ at $m(\bar\theta)$; in interior–regime formulas we evaluate at $m(\theta^*)$.

IA.3.2 Working assumptions for comparative statics

All comparative-statics results below are local and use the following assumptions:

(A1) **Local shocks & no regime switch.** Changes in primitives are infinitesimal (so derivatives are informative) and do not switch the regime (cap: $\theta=\bar\theta$; interior: $0<\theta^*<\bar\theta$).

(A2) **Pledge frontier.** $m'(\cdot)>0$, $m''(\cdot)<0$.

(A3) **Advance rate.** $R>\pi, 0<\rho<1$ so that $y(R,\pi,\rho)$ is well-defined and $\chi_R = \partial \ln y/\partial \ln R = -R/(R-\pi) < 0$.

(A4) **Positivity & feasibility.** $1 + \frac{h}{p_k} - y\, m(\theta) > 0$ for $\theta \in [0, \bar{\theta}]$, ensuring $M > 0$ and a nonempty feasible set.

(A5) **Main-region condition at the cap (when invoked).** For cap–regime h, τ shocks we use $\kappa = \frac{\tau c}{2M^2} < 1$, which holds empirically at observed scales and under which the exposure effect dominates the brake.

(A6) **Compact neighborhoods.** Primitives vary over compact neighborhoods where the objects above remain well-defined.

(A7) **Incomplete pledge adjustment** In the interior region ($0 < \theta^* < \bar{\theta}$), a firm cannot keep buying the same number of robots after a price or rate increase purely by borrowing more; holding some extra own cash per robot is always required: $0 \le W_\theta \eta_{\theta,p} < W_{p_k}, \qquad W_\theta \eta_{\theta,R} + W_f \chi_R < 0.$

IA.3.3 Competitive pricing, objective neutrality of b' and the b'–FOC

This subsection derives the condition $\lambda = \mu$ used in the main text.

Step 1: Zero–NPV financing and reduced–form objective. Under competitive, risk–neutral pricing, the loan is priced so that the lender's expected discounted payoff equals the amount lent:

$$\mathbb{E}_0\left[\frac{\text{repayment}}{R}\right] = \frac{\pi\, b' + (1-\pi)\, m(\theta)\, \rho\, p_k k'}{R} = b'.$$

From the firm's perspective, the present value of the financing leg is therefore

$$\text{PV(financing)} = \underbrace{b'}_{\text{inflow at } t=0} \; \underbrace{- \mathbb{E}_0\left[\tfrac{\text{repayment}}{R}\right]}_{\text{outflow in PV}} = b' - b' = 0.$$

Hence issuing or repaying one more unit of debt has zero net present value in the objective; b' affects value only through the constraints.

Using this pricing condition, the firm's objective can be written purely in terms of real cash flows, as in (3.6):

$$\mathcal{U}(k', \theta) = \frac{\pi}{R}\Big[A(k')^\alpha + (1-\delta)p_k k'\Big] + \frac{1-\pi}{R}\Big[\rho(1-\theta)p_k k'\Big] - (p_k + h)k' - \tfrac{\tau}{2}(k')^2.$$

Step 2: Lagrangian with budget and capacity constraints. The choice variables are (k', b', θ) and the constraints are:

- Budget identity (3.5):
$$c + b' = (p_k + h)k' + \tfrac{\tau}{2}(k')^2.$$
- Borrowing capacity (3.3):
$$b' \le y(R, \pi, \rho)\, m(\theta)\, p_k k'.$$
- Pledge cap (3.4):
$$0 \le \theta \le \bar{\theta}.$$

Let $\lambda \ge 0$ be the multiplier on (3.5), $\mu \ge 0$ the multiplier on the borrowing–capacity constraint (3.3), and $\nu^-, \nu^+ \ge 0$ the multipliers on the lower and upper bounds for θ. Writing the borrowing constraint in the standard $g(\cdot) \le 0$ form as $y\, m(\theta)\, p_k k' - b' \ge 0$, the Lagrangian is

$$\begin{aligned}\mathcal{L} = \mathcal{U}(k', \theta) &+ \lambda\Big(c + b' - (p_k + h)k' - \tfrac{\tau}{2}(k')^2\Big) \\ &+ \mu\Big(b' - y(R, \pi, \rho)\, m(\theta)\, p_k k'\Big) + \nu^-\,(0 - \theta) + \nu^+\,(\theta - \bar{\theta}).\end{aligned}$$

Step 3: b'–FOC and equality of multipliers. Crucially, after using the zero–NPV condition, b' appears *only* in the constraints inside the Lagrangian. Differentiating with respect to b' gives

$$\frac{\partial \mathcal{L}}{\partial b'} = \lambda \underbrace{\frac{\partial}{\partial b'}\Big(c + b' - (p_k + h)k' - \tfrac{\tau}{2}(k')^2\Big)}_{=1} + \mu \underbrace{\frac{\partial}{\partial b'}\Big(y\, m(\theta)\, p_k k' - b'\Big)}_{=-1} = \lambda - \mu.$$

The first–order condition in an interior optimum is

$$\frac{\partial \mathcal{L}}{\partial b'} = 0 \quad \Longrightarrow \quad \boxed{\lambda = \mu}.$$

Interpretation. The equality $\lambda = \mu$ reflects the neutrality of the loan under competitive pricing. An extra dollar of b' relaxes the budget constraint by one dollar at $t = 0$ (coefficient $+1$ in the budget constraint) but uses up one dollar of borrowing capacity (coefficient -1 in the collateral constraint). Because the financing leg itself has zero net present value, these are the only channels through which b' affects value, and the b'–FOC simply equates the shadow values:

λ = marginal value of one more dollar of internal funds = marginal value of one more dollar of borrowing capacity = μ.

When finance is slack, both constraints are slack so $\lambda = \mu = 0$; when finance binds, we have $\lambda = \mu > 0$ and we use the common multiplier λ as our summary measure of tightness in the binding region.

IA.3.4 Invoice sufficiency lemma

This subsection states and proves a basic lemma: when the borrowing constraint binds, the scale k' is determined entirely by the invoice M via the cash identity $k'M = c$, so local elasticities of k' with respect to any primitive x are given by $-\partial \ln M / \partial \ln x$; when finance is slack, k' is determined instead by the technology FOC and depends on $A(\cdot)$.

Lemma 1 (Invoice sufficiency)**.** *(i) For any feasible (k', θ) the invoice $M(k', \theta)$ in (IA.3.1) is well-defined.*

(ii) If finance binds in equilibrium, so that the borrowing constraint holds with equality at some θ^*, then k' is pinned by $k'M(k', \theta^*) = c$, and for any primitive $x \in \{p_k, R, h, \pi, \rho\}$,

$$\frac{d \ln k'}{d \ln x} \;=\; -\,\frac{d \ln M(k', \theta^*; x)}{d \ln x},$$

with all dependence on the production function $A(\cdot)$ absorbed into the fact that it selects the binding region.

(iii) If finance is slack (so $\lambda = \mu = 0$ and $\theta^* = 0$), then k' is determined by the unconstrained technology FOC and the elasticities $\partial \ln k' / \partial \ln x$ depend on $A'(\cdot)$ and $A''(\cdot)$; in this case M is not a sufficient statistic for elasticities.

Proof. (i) For any given (k', θ) satisfying the feasibility conditions in Section IA.3.1, $M(k', \theta)$ is defined by (IA.3.1); it is simply the own–cash outlay per unit of k'.

(ii) When finance binds, the borrowing constraint holds with equality at θ^*: $b' = y\, m(\theta^*) p_k k'$. Substituting into the budget identity,

$$c + b' = (p_k + h)k' + \tfrac{\tau}{2}(k')^2 \quad \Longrightarrow \quad k'M(k', \theta^*) = c.$$

Taking logs and differentiating at fixed c gives

$$d \ln k' + d \ln M(k', \theta^*; x) = 0 \quad \Rightarrow \quad \frac{\partial \ln k'}{\partial \ln x} = -\,\frac{\partial \ln M(k', \theta^*; x)}{\partial \ln x}.$$

The dependence of θ^* on x is handled by its own FOC/KKT conditions; the production function only determines whether the solution lies in the binding region and does not appear in $\partial \ln M / \partial \ln x$.

(iii) When finance is slack, the borrowing constraint and cap are not binding, $\lambda = \mu = 0$ and $\theta^* = 0$, and the optimal scale k' is determined instead by the unconstrained FOC

$$\frac{\pi}{R}\,\alpha\,A(k')^{\alpha-1} = \underbrace{(p_k + h) + \tau k'}_{\text{gross outlay}} - \underbrace{\left[\tfrac{\pi}{R}(1-\delta) + \tfrac{1-\pi}{R}\,\rho\right] p_k}_{\text{expected resale}}.$$

By the implicit function theorem, $\partial \ln k' / \partial \ln x$ depends on $A'(\cdot), A''(\cdot)$ through the derivative of the left-hand side. In this region, k' is not determined by $k'M = c$, so the mapping $x \mapsto M$ is not sufficient to recover $\partial \ln k' / \partial \ln x$. □

IA.3.5 Cap–regime derivations

This section provides the closed-form elasticities used in the main text for the cap regime ($\theta = \bar{\theta}$), where the pledge cap binds and k' is pinned by $k'M = c$.

Evaluation. Throughout this subsection, $m(\cdot)$ is evaluated at $m(\bar{\theta})$.

IA.3.5.1 Price elasticity at the cap

With $\theta = \bar{\theta}$, (IA.3.1) becomes

$$k'\,M = c, \qquad M = p_k\Big(1 + \frac{h}{p_k} - y\,m(\bar{\theta})\Big) + \frac{\tau}{2}k'.$$

Hence

$$\frac{\partial \ln M}{\partial \ln p_k} = W_{p_k} \equiv \frac{p_k\big(1 - y\,m(\bar{\theta})\big)}{M}, \qquad \frac{\partial \ln M}{\partial \ln k'} = \kappa \equiv \frac{\frac{\tau}{2}k'}{M}.$$

Using $d\ln k' + d\ln M = 0$,

$$(1+\kappa) d\ln k' = -W_{p_k}\, d\ln p_k \quad \Longrightarrow \quad \boxed{\varepsilon_{k',p_k} = -\frac{W_{p_k}}{1+\kappa} < 0}.$$

IA.3.5.2 Rate elasticity at the cap

At the cap,

$$\frac{\partial \ln M}{\partial \ln R} = \frac{\partial \ln M}{\partial \ln y} \cdot \frac{\partial \ln y}{\partial \ln R} = -\,W_f\,\chi_R, \qquad W_f = \frac{p_k\,y\,m(\bar{\theta})}{M}, \quad \chi_R < 0.$$

Then

$$(1+\kappa) d\ln k' = -\big(-\,W_f\,\chi_R\big)\, d\ln R \quad \Longrightarrow \quad \boxed{\varepsilon_{k',R} = \frac{W_f\,\chi_R}{1+\kappa} < 0}.$$

IA.3.6 Interior–regime derivations

This section provides the closed-form elasticities used in the main text for the interior regime ($0 < \theta^* < \bar{\theta}$), where pledge adjusts at the margin.

Evaluation. Throughout this subsection, $m(\cdot)$ is evaluated at $m(\theta^*)$.

IA.3.6.1 Pledge elasticities: definition and origin

In the interior regime ($0 < \theta^* < \bar{\theta}$) the optimal pledge is endogenous and moves with shocks. We define the *pledge elasticities*

$$\eta_{\theta,p} \equiv \frac{\partial \ln \theta^*}{\partial \ln p_k}, \qquad \eta_{\theta,R} \equiv \frac{\partial \ln \theta^*}{\partial \ln R}.$$

They are obtained by differentiating the interior first–order condition for pledge,

$$\lambda\, y\, m'(\theta^*) = \frac{1-\pi}{R}\,\rho, \tag{IA.3.4}$$

which holds when the borrowing constraint binds and θ is interior.

Price shock. Holding (R, π, ρ) fixed (so y is fixed), log–differentiating (IA.3.4) with respect to p_k gives

$$d\ln\lambda + d\ln m'(\theta^*) = 0, \qquad d\ln m'(\theta^*) = \frac{\theta^* m''(\theta^*)}{m'(\theta^*)}\, d\ln\theta^*.$$

Hence

$$\eta_{\theta,p} = \frac{\partial \ln \theta^*}{\partial \ln p_k} = -\,\frac{m'(\theta^*)}{\theta^* m''(\theta^*)}\,\frac{\partial \ln \lambda}{\partial \ln p_k}. \tag{IA.3.5}$$

Under $m''(\cdot) < 0$ and the mild monotonicity $\partial \ln\lambda / \partial \ln p_k > 0$ (a higher p_k tightens the invoice M and raises the shadow value of funds), we have $\eta_{\theta,p} > 0$.

Rate shock. Log–differentiating (IA.3.4) with respect to R yields

$$\frac{\partial \ln \lambda}{\partial \ln R} + \chi_R + \frac{\theta^* m''(\theta^*)}{m'(\theta^*)}\,\eta_{\theta,R} = -1,$$

so

$$\eta_{\theta,R} = \Big(-1 - \frac{\partial \ln \lambda}{\partial \ln R} - \chi_R\Big)\,\frac{m'(\theta^*)}{\theta^* m''(\theta^*)}. \tag{IA.3.6}$$

Without additional structure this elasticity is not signed from primitives alone; in what follows we rely on Assumption (A7), which directly restricts the combination $W_\theta \eta_{\theta,R} + W_f \chi_R$ that enters the rate elasticity.

IA.3.6.2 Price elasticity with interior pledge

With c fixed and $d\ln\theta = \eta_{\theta,p}\, d\ln p_k$,

$$d\ln M = W_{p_k}\, d\ln p_k + \kappa\, d\ln k' - W_\theta\, d\ln\theta = W_{p_k}\, d\ln p_k + \kappa\, d\ln k' - W_\theta\, \eta_{\theta,p}\, d\ln p_k.$$

Using $d\ln k' + d\ln M = 0$,

$$(1+\kappa)\, d\ln k' = -\big(W_{p_k} - W_\theta\, \eta_{\theta,p}\big)\, d\ln p_k \quad \Longrightarrow \quad \boxed{\varepsilon^{\text{int}}_{k',p_k} = -\frac{W_{p_k} - W_\theta\, \eta_{\theta,p}}{1+\kappa}}.$$

Concavity of m and the interior θ–FOC (see §IA.3.7) imply

$$0 \le W_\theta\, \eta_{\theta,p} < W_{p_k}.$$

IA.3.6.3 Rate elasticity with interior pledge

Using $\partial \ln M / \partial \ln R = -W_f\, \chi_R$ and $\partial \ln M / \partial \ln \theta = -W_\theta$, and setting $d\ln\theta = \eta_{\theta,R}\, d\ln R$, we obtain

$$d\ln M = \big(-W_f\, \chi_R - W_\theta\, \eta_{\theta,R}\big)\, d\ln R + \kappa\, d\ln k'.$$

Combining this with $d\ln k' + d\ln M = 0$ yields

$$(1+\kappa)\, d\ln k' = \big(W_f\, \chi_R + W_\theta\, \eta_{\theta,R}\big)\, d\ln R,$$

so that

$$\varepsilon^{\text{int}}_{k',R} \equiv \frac{\partial \ln k'}{\partial \ln R} = \frac{W_f\, \chi_R + W_\theta\, \eta_{\theta,R}}{1+\kappa}. \tag{IA.3.7}$$

Under Assumption (A7), $W_\theta \eta_{\theta,R} + W_f \chi_R < 0$, so the numerator in (IA.3.7) is negative and hence $\varepsilon^{\text{int}}_{k',R} < 0$.

IA.3.7 Sign preservation in the interior regime (Proposition 1)

This section proves Proposition 1 in the main text.

With binding finance and $0 < \theta^* < \bar\theta$, the θ–FOC is

$$\lambda\, y\, m'(\theta^*) = \frac{1-\pi}{R}\rho.$$

Taking logs and differentiating:

(i) Price shock. From (IA.3.5), $m'' < 0$ and the mild monotonicity $\partial \ln\lambda / \partial \ln p_k > 0$ imply $\eta_{\theta,p} > 0$: a higher sticker price makes encumbrance more attractive. Assumption (A7) states that in the interior region pledge adjustment cannot fully offset this higher price: $0 \le W_\theta \eta_{\theta,p} < W_{p_k}$. Since $W_{p_k} > 0$, this implies $W_{p_k} - W_\theta \eta_{\theta,p} > 0$, and from (IA.3.6.2) we obtain $\varepsilon^{\text{int}}_{k',p_k} < 0$.

(ii) Rate shock. Differentiating $\ln\lambda + \ln y + \ln m'(\theta^*) = \ln(1-\pi) + \ln\rho - \ln R$ with respect to $\ln R$ yields (IA.3.6). As noted in §IA.3.6.1, this does not by itself sign $\eta_{\theta,R}$. However, Assumption (A7) requires that pledge adjustment cannot fully undo a tightening in financing conditions: $W_\theta \eta_{\theta,R} + W_f \chi_R < 0$. Together with (IA.3.7), this implies $\varepsilon^{\text{int}}_{k',R} < 0$.

Under Assumptions (A1)–(A7), we therefore have $\varepsilon^{\text{int}}_{k',p_k} < 0$ and $\varepsilon^{\text{int}}_{k',R} < 0$, which proves Proposition 1.

IA.3.8 Lemmas on the shadow value λ

Let $\lambda \ge 0$ be the multiplier on (IA.3.1) (and μ the multiplier on capacity); by §IA.3.3, $\lambda = \mu$ when finance binds.

Lemma 2 (Interior). *If $0 < \theta^* < \bar\theta$, then*

$$\lambda = \frac{(1-\pi)\rho}{R\, y\, m'(\theta^*)}, \qquad \frac{\partial\lambda}{\partial\theta^*} = -\frac{(1-\pi)\rho}{R\, y}\frac{m''(\theta^*)}{(m'(\theta^*))^2} > 0.$$

Lemma 3 (Cap). *If $\theta = \bar\theta$ binds, then $\lambda \ge \dfrac{(1-\pi)\rho}{R\, y\, m'(\bar\theta)}$.*

IA.3.9 Region selection (slack vs. binding)

This lemma is used in the "Role of Technology" discussion in the main text.

Let k^{fb} solve the slack FOC with $\theta = 0$ (MPK = user cost net of expected resale). Finance binds iff the own–cash requirement at k^{fb} exceeds available cash even at the pledge cap:

$$c \;<\; \min_{\theta\in[0,\bar\theta]} \; k^{\text{fb}}\Big[p_k\Big(1 + \frac{h}{p_k} - y\, m(\theta)\Big) + \frac{\tau}{2}\, k^{\text{fb}}\Big] \;=\; k^{\text{fb}}\Big[p_k\Big(1 + \frac{h}{p_k} - y\, m(\bar\theta)\Big) + \frac{\tau}{2}\, k^{\text{fb}}\Big].$$

IA.3.10 Slack (cash–rich) benchmark: FOC and elasticities

This section underlies the slack (cash–rich) benchmark in the main text.

With $\theta = 0$ and slack finance,

$$\frac{\pi}{R}\,\alpha A(k')^{\alpha-1} = (p_k + h) + \tau k' - \left[\tfrac{\pi}{R}(1-\delta) + \tfrac{1-\pi}{R}\rho\right] p_k.$$

Let $F(k', p_k, h; R, \pi, \rho) = 0$ denote this equation. By the implicit–function theorem,

$$\frac{\partial \ln k'}{\partial \ln p_k} = \frac{\frac{p_k}{k'}\left(1 - \frac{\pi}{R}(1-\delta) - \frac{1-\pi}{R}\rho\right)}{\frac{\pi}{R}\,\alpha(\alpha-1)\,A(k')^{\alpha-2} - \tau}, \qquad \frac{\partial \ln k'}{\partial \ln h} = \frac{\frac{h}{k'}}{\frac{\pi}{R}\,\alpha(\alpha-1)\,A(k')^{\alpha-2} - \tau},$$

and rate/(π, ρ) elasticities follow analogously (both sides move by discounting/resale).

IA.3.11 Proof of Proposition 2 (second–order cross–effects)

Assume small shocks in $x \in \{p_k, R\}$ do not switch regimes (interior $0 < \theta^* < \bar{\theta}$ or cap $\theta^* = \bar{\theta}$). In the binding region, $E_x \equiv \partial \ln k'/\partial \ln x$ satisfies

$$\frac{\partial E_x}{\partial z} = -\frac{\partial}{\partial z}\left(\frac{\partial \ln M}{\partial \ln x}\right) = -\frac{\partial^2 \ln M}{\partial \ln x\, \partial z},$$

that is, for $z > 0$,

$$\boxed{\frac{\partial E_x}{\partial z} = -\frac{1}{z}\,\frac{\partial^2 \ln M}{\partial \ln x\, \partial \ln z}}.$$

Proof. From $\ln k' + \ln M(k', \theta^*; x, z) = \ln c$ we have $E_x = -\partial \ln M/\partial \ln x$. Differentiating w.r.t. z along the equilibrium manifold $(k'(x,z), \theta^*(x,z))$ yields the first line; the log–level rescaling is immediate since $\partial/\partial z = (1/z)\,\partial/\partial \ln z$ for $z > 0$. □

IA.3.12 Notes for Tables 1 and 2 (arrows and diagnostics)

This section derives the sign arrows for λ, θ^*, and the price and rate elasticities reported in Tables 1 and 2 in the main text. The first part treats the cap regime; the second part treats the interior regime.

Assumption MR (main region at the cap). When we evaluate the impact of h and τ on elasticities at the cap, we use the "main region" condition

$$\kappa \equiv \frac{\tau c}{2M^2} < 1,$$

which is equivalent to $M^2 > B$ with $B \equiv \frac{\tau c}{2}$ in the closed forms below. Empirically, at observed scales κ is modest, and this assumption delivers the sign patterns reported in Table 1.

Closed forms at the cap. Write

$$|\,\varepsilon_{k',p_k}\,| = \frac{W_{p_k}}{1+\kappa} = \frac{A/M}{1 + B/M^2} = \frac{A\,M}{M^2 + B}, \qquad |\,\varepsilon_{k',R}\,| = \frac{W_f\,|\chi_R|}{1+\kappa} = \frac{C/M}{1 + B/M^2} = \frac{C\,|\chi_R|\,M}{M^2 + B},$$

with $A \equiv p_k(1 - y\,m(\bar{\theta}))$, $C \equiv p_k\,y\,m(\bar{\theta})$, and $B \equiv \frac{\tau c}{2}$.

At the cap.

- $c\downarrow$**.** Holding M fixed to first order in the direct channel (the induced change in M via k' is second order

under a small shock),

$$\frac{\partial|\,\varepsilon_{k',p_k}\,|}{\partial c}=\frac{\partial}{\partial c}\Big(\frac{W_{p_k}}{1+\kappa}\Big)=-\frac{W_{p_k}}{(1+\kappa)^2}\,\frac{\partial\kappa}{\partial c}=-\frac{W_{p_k}}{(1+\kappa)^2}\,\frac{\tau}{2M^2}<0,$$

and the same sign obtains for $|\,\varepsilon_{k',R}\,|$. Thus both magnitudes rise when c falls; additionally, $\lambda\uparrow$ (tighter budget).

- $h\downarrow$**.** Then $M\downarrow$. Using the closed form,

$$\frac{\partial|\,\varepsilon_{k',p_k}\,|}{\partial M}=A\,\frac{B-M^2}{(M^2+B)^2}\;<\;0\quad\text{under MR }(\kappa<1\Leftrightarrow M^2>B),$$

so with $dM<0$ we have $d|\,\varepsilon_{k',p_k}\,|>0$. An identical calculation yields $d|\,\varepsilon_{k',R}\,|>0$. Meanwhile $\lambda\downarrow$ (easier finance).

- $\tau\downarrow$**.** Here $B\downarrow$ and

$$\frac{\partial|\,\varepsilon_{k',p_k}\,|}{\partial B}=-\,\frac{A\,M}{(M^2+B)^2}<0,\qquad\frac{\partial|\,\varepsilon_{k',R}\,|}{\partial B}=-\,\frac{C\,|\chi_R|\,M}{(M^2+B)^2}<0,$$

so both magnitudes rise; $\lambda\downarrow$ (easier finance).

- $\bar\theta\downarrow$**.** Then $m(\bar\theta)\downarrow\Rightarrow A\uparrow, C\downarrow$. Hence $|\varepsilon_{k',p_k}|\uparrow$ while $|\varepsilon_{k',R}|\downarrow$; also $\lambda\uparrow$ (tighter cap).
- $y\uparrow$**.** Then $A\downarrow, C\uparrow$. Hence $|\varepsilon_{k',p_k}|\downarrow$, $|\varepsilon_{k',R}|\uparrow$; and $\lambda\downarrow$ (safer debt).

Interior (bounds and arrows). Recall

$$\varepsilon^{\text{int}}_{k',p_k}=-\frac{W_{p_k}-W_\theta\,\eta_{\theta,p}}{1+\kappa},\quad\varepsilon^{\text{int}}_{k',R}=\frac{W_f\,\chi_R+W_\theta\,\eta_{\theta,R}}{1+\kappa}.$$

Therefore, using (A1)–A(7):

- $c\downarrow$: $\theta^*\uparrow$ and both magnitudes rise (exposure rises; cushion cannot fully offset); $\lambda\uparrow$.
- $h\downarrow$: $\theta^*\downarrow$ while both magnitudes rise (exposure up; cushion shrinks); $\lambda\downarrow$.
- $\tau\downarrow$: $\kappa\downarrow$ and both magnitudes rise; θ^* nearly flat to first order.
- $y\uparrow$: price sensitivity falls, rate sensitivity rises; $\theta^*\downarrow$.
- $\bar\theta$ non–binding: no first–order effect on θ^*, elasticities, or λ.

IA.3.13 Proof of Proposition 3 (diagnostic uniqueness)

We now formalize the "diagnostic" result used in the text. This proposition provides the formal version of the informal statement in Proposition 3 in the main text.

Proposition 4 (Diagnostic uniqueness of liquidity)**.** *Fix a regime (cap: $\theta=\bar\theta$; interior: $0<\theta^*<\bar\theta$) and consider an infinitesimal change in a single primitive $z\in\{c,h,\tau,y,\bar\theta\}$ (cap) or $z\in\{c,h,\tau,y\}$ (interior) that does not switch the regime. Assume $m'(\cdot)>0$, $m''(\cdot)<0$, $\chi_R<0$, the positivity condition in §IA.3.1, and, at the cap for h,τ shocks, Assumption MR ($\kappa<1$).*

Cap. *Define* $S_{\text{cap}}(z)\equiv\big(\operatorname{sign}\Delta\lambda,\operatorname{sign}\Delta|\varepsilon_{k',p_k}|,\operatorname{sign}\Delta|\varepsilon_{k',R}|\big)$. *Then*

$$S_{\text{cap}}(c\downarrow)=(+,+,+),$$

and for each $z \in \{h\downarrow, \tau\downarrow, \bar{\theta}\downarrow, y\uparrow\}$ *the triple* $S_{\text{cap}}(z)$ *differs from* $(+,+,+)$.

Interior. *Define* $S_{\text{int}}(z) \equiv \big(\operatorname{sign}\Delta\theta^*, \operatorname{sign}\Delta|\varepsilon_{k',p_k}|, \operatorname{sign}\Delta|\varepsilon_{k',R}|\big)$. *Then*

$$S_{\text{int}}(c\downarrow) = (+,+,+),$$

and for each $z \in \{h\downarrow, \tau\downarrow, y\uparrow\}$ *the triple* $S_{\text{int}}(z)$ *differs from* $(+,+,+)$ *(if* $\bar{\theta}$ *is slack, it has no first-order effect).*

Proof. Cap. The derivatives above show (i) $c\downarrow$ raises λ and both magnitudes; (ii) $h\downarrow$ or $\tau\downarrow$ raise both magnitudes but lower λ (MR used for the magnitudes); (iii) $\bar{\theta}\downarrow$ raises price sensitivity, lowers rate sensitivity, and raises λ; (iv) $y\uparrow$ raises rate sensitivity, lowers price sensitivity, and lowers λ. Hence only $c\downarrow$ yields $(+,+,+)$.

Interior. Using (A1)-A(7): (i) $c\downarrow$ raises θ^* and both magnitudes; (ii) $h\downarrow$ lowers θ^* while raising both magnitudes; (iii) $\tau\downarrow$ leaves θ^* approximately flat while raising both magnitudes; (iv) $y\uparrow$ lowers price sensitivity and raises rate sensitivity, with $\theta^*\downarrow$. Thus only $c\downarrow$ yields $(+,+,+)$. □

IA.3.14 Asset-Based vs. Cash-Flow-Based Lending and the Borrowing Constraint

Our model writes the borrowing constraint of a given firm as

$$b' \;\le\; y\, m(\theta)\, p_k\, k', \tag{IA.3.8}$$

so that, under binding finance, the scale of robot investment is pinned by the cash identity

$$c \;=\; k'\, M(k',\theta), \qquad M \;=\; p_k\Big(1 + a - y\, m(\theta)\Big) + \frac{\tau}{2}k', \tag{IA.3.9}$$

with $a \equiv h/p_k$. A natural question, in light of Lian and Ma (2021), is how (IA.3.8) relates to the empirical distinction between "asset-based" lending (ABL) and "cash-flow-based" lending. This subsection clarifies two points:

(i) how both ABL and earnings-based lending can be mapped into $(y, m(\cdot))$; and

(ii) why our comparative statics are correctly interpreted as conditional on $(c, h, \tau, y, m(\cdot))$, regardless of whether the underlying contract is classified as asset-based or cash-flow-based.

(i) Asset-based lending. In the simplest ABL implementation the lender sets a constant advance rate $\phi \in (0,1)$ on the liquidation value of robots:

$$b' \;\le\; \phi\, \rho\, p_k\, k'.$$

Comparing with (IA.3.8), this corresponds to choosing $y\, m(\theta) = \phi\, \rho$, i.e. the per-unit financeable share of the purchase price is a constant fraction of resale value. In this case the primitives y and $m(\theta)$ literally reflect physical collateral and haircuts; the cash identity (IA.3.9) simply subtracts this financed share from the per-unit outlay and adds the installation term.

(ii) Cash-flow-based (earnings-based) lending. For large US firms, Lian and Ma (2021) document that the majority of debt by value is governed by earnings-based covenants rather than by classic ABL. A typical covenant has the form

$$\text{Debt}_i \;\le\; \kappa_i \times \text{EBITDA}_i, \tag{IA.3.10}$$

possibly combined with interest-coverage constraints. Let firm i have baseline EBITDA E_{0i} and incremental EBITDA per robot $\Delta E_i(k')$; after investing k' robots, EBITDA is

$$\text{EBITDA}_i(k') = E_{0i} + \Delta E_i(k').$$

Substituting into (IA.3.10),

$$\text{Debt}_i(k') \;\le\; \kappa_i\big[E_{0i} + \Delta E_i(k')\big].$$

We do not model the joint determination of all long-term and short-term debt for the firm. Instead, we take the robot project as incremental and ask how much *additional* borrowing capacity one more unit of k' generates at the margin, given the lender's internal rule. A first-order expansion around the current scale $\bar{k}'$ yields

$$\Delta\text{Debt}_i \;\approx\; \kappa_i\,\Delta E_i'(\bar{k}')\,\Delta k'.$$

Define the *incremental borrowing capacity per unit of robot spending* as

$$\Phi_i \;\equiv\; \frac{\partial \text{Debt}_i}{\partial k'}\Big|_{k'=\bar{k}'}\Big/ p_k \;=\; \frac{\kappa_i\,\Delta E_i'(\bar{k}')}{p_k}.$$

For small changes in k' around $\bar{k}'$, the earnings-based constraint is therefore locally equivalent to

$$b_i' \;\le\; \Phi_i\, p_k\, k_i',$$

which is of exactly the same form as (IA.3.8) with $y_i m_i(\theta_i) \equiv \Phi_i$. In words, Φ_i summarizes how many dollars of incremental debt the lender is willing to extend per dollar of robot expenditure, taking into account all of the details of firm i's technology, baseline earnings, and covenant structure.

The key point is that, for the small shocks we study (changes in p_k and R around the pre-shock scale), both ABL and cash-flow-based contracts can be expressed in the reduced form (IA.3.8) with firm-specific $(y_i, m_i(\cdot))$ that encode the *effective financeable share* of robot investment for firm i.

(iii) Conditional comparative statics. Once we summarize the lender–firm relationship for each firm i by its financing primitives $(y_i, m_i(\cdot))$, all of our comparative statics are derived *conditional* on the tuple $(c_i, h_i, \tau_i, y_i, m_i(\cdot))$. In the binding region, scale is pinned by the cash identity

$$c_i = k_i' M_i(k_i', \theta_i),$$

and the elasticities of robot investment with respect to the price of capital and the policy rate are

$$\varepsilon^{(i)}_{k',p_k} = -\frac{W^{(i)}_{p_k}}{1+\kappa_i}, \qquad \varepsilon^{(i)}_{k',R} = -\frac{W^{(i)}_{f}\,\chi_R}{1+\kappa_i}$$

at the cap, and analogous expressions in the interior regime, where $W^{(i)}_{p_k}, W^{(i)}_{f}, \kappa_i$ are functions of $(c_i, h_i, \tau_i, y_i, m_i(\cdot))$. Crucially, these formulas do *not* depend on how y_i and $m_i(\cdot)$ are generated—whether through explicit liens on robots, earnings-based covenants, or a mixture of both. Differences in firms' technologies and business models that affect banks' willingness to lend enter only through the induced values of $(y_i, m_i(\cdot))$; given those primitives, the small-shock responses of k_i' to p_k and R are exactly as in our model.

In this sense, our theory should be read as providing comparative statics *conditional on effective financing capacity*, consistent with the evidence in Lian and Ma (2021) and Drechsel (2023) that large-firm borrowing limits are often earnings-based rather than purely collateral-based. Some firms in our sample will reach their borrowing frontier primarily through asset-based channels, others through earnings-based covenants, but for both types the relevant primitives for our analysis are the implied $(c_i, h_i, \tau_i, y_i, m_i(\cdot))$ summarizing how much of the robot "invoice" can be financed externally and how much must be paid out of internal cash.