Who Adopts New Technologies?

Financing Constraints, Intangible Investment, and the Diffusion of Automation*

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

This version: July 10, 2026

Abstract:

We study how financial constraints shape automation investment using transaction-level import records of robots and automation machinery linked to Compustat manufacturers. Exploiting exogenous shocks to automation prices (the yen depreciation) and financing costs (monetary policy surprises), we find that financially constrained firms are roughly twice as responsive to both—contrary to canonical models in which financing frictions mute investment responses. The constrained–unconstrained amplification is 56 percent larger for the price shock than for the rate shock. A collateral model explains the pattern: new technologies like robots bundle pledgeable hardware with non-pledgeable intangibles that must be paid from internal cash, so a fall in hardware prices relaxes the binding cash constraint disproportionately for constrained firms. Under the model's benchmark binding-regime and lending-pass-through assumptions, the estimates imply hardware advance rates near one-half and about ten percentage points lower for constrained firms, broadly consistent with debt-financing patterns among adopters. The results highlight how the financing of non-leverageable setup costs, rather than of the equipment itself, shapes which firms adopt new technologies.

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

*We are grateful for valuable comments from Martin Aragoneses, Murray Frank, Jordan Richmond and seminar participants at INSEAD, Imperial College, SKEMA Business School, the Virtual Corporate Finance Seminar, the University of Maryland Brown Bag, the Smith School FEST seminar series, GWU-CIBER conference on Geopolitics and Innovation, the Georgetown Center for Economic Research alumni conference, the International Finance Corporation, the North American Summer Meeting of the Econometric Society (NASMES), the Washington Area International Trade Symposium (WAITS), the 9th Center for Corporate Finance Conference at Norwegian School of Economics, and the March 2026 MFA meetings. Any mistakes are our own.

†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

Between 2012 and 2015, a sharp yen depreciation lowered the dollar price of Japanese machinery. Yet the firms that increased automation most were those with the least financial flexibility. What determines which firms adopt new technologies? In canonical investment models in which financing constraints restrict the ability to scale a fixed opportunity set, constrained firms respond less to favorable user-cost shocks. We document the opposite. When automation becomes cheaper or easier to finance, financially constrained firms are the marginal adopters, responding more strongly than unconstrained firms to both declines in capital-goods prices and financing costs. This reversal suggests that the financing friction governing technology adoption differs fundamentally from those underlying standard investment models, which study firm responses to changes in discount rates, internal funds, or demand.[1] Much less is known about how financing constraints interact with the price of the capital good itself, particularly for technologies requiring both tangible and intangible investment.

Our central insight is that new technologies combine two very different investments: tangible hardware that lenders can finance against collateral and complementary, largely non-pledgeable intangible investments—software integration, customization, worker training, and organizational change—that must be financed internally. The financing friction governing technology adoption arises from this combination. The relevant financing constraint is therefore not the machine's purchase price alone, but the cash required to make it productive. When the hardware becomes cheaper or borrowing conditions improve, that cash requirement falls disproportionately for firms operating near their liquidity constraint, making financially constrained firms the marginal adopters.

Industrial robots provide an unusually clean setting to study this mechanism. Their hardware is standardized and pledgeable, while complementary integration and setup costs are substantial and well documented. Unlike most capital goods, adoption is observed directly, purchase by purchase. We construct a transaction-level measure of automation adoption for U.S. public manufacturers by linking bill-of-lading import records to Compustat between 2007 and 2019, allowing us to observe the

[1] See Bernanke and Gertler (1995); Gilchrist and Zakrajšek (2012); Ottonello and Winberry (2020); Gopinath et al. (2017) and the investment–cash-flow literature, e.g., Fazzari et al. (1988); Whited (1992); Hennessy et al. (2007); Almeida and Campello (2007).

timing, type, and origin of individual automation purchases directly rather than infer adoption from capital expenditures or industry-level robot counts. Relative to sources such as the International Federation of Robotics (IFR), these data offer substantially greater granularity and temporal precision. Together, these features allow us to separately identify the two components of automation's user cost—its capital-goods price and its financing cost—for the same investment decision, something previous work has largely been unable to do.

Our empirical design exploits two plausibly exogenous shocks that separately shift the two components of the user cost of automation: its capital-goods price and its financing cost.

The first exploits the structure of global robot supply to generate price variation orthogonal to U.S. firm fundamentals. Production is concentrated among a handful of manufacturers outside the U.S., the relevant import categories are narrowly identified by HS codes, and Japan alone accounts for more than half of U.S. industrial robot imports—so yen-dollar movements translate directly into dollar-price declines for U.S. buyers (the import price index for Japanese machinery fell more than 10 percent against the depreciation). Critically, the 2012-2015 depreciation was driven by Japanese domestic monetary policy (Abenomics) rather than U.S. macroeconomic conditions or sector-specific demand, and the concentrated supply structure left U.S. firms limited scope to substitute toward non-Japanese suppliers at the intensive margin.

Our second shock isolates variation in financing costs. Following Bauer and Swanson (2023), we instrument changes in the one-year Treasury yield using orthogonalized high-frequency monetary-policy surprises that purge Fed-information effects documented in prior work. Because both shocks operate on the same investment margin, our setting allows us to compare how the same firms respond to changes in capital-goods prices versus financing costs.

We find that financially constrained firms respond roughly twice as strongly as unconstrained firms to both shocks. Quantitatively, a 10% yen depreciation raises the unconditional probability that a U.S. manufacturing firm begins importing robots by 2.25 percentage points; the effect is 1.4 to 1.8 percentage points larger among firms classified ex ante as financially constrained — by below-median cash holdings, by the debt-constraint language identified in 10-K filings (**?**Linn and Weagley, 2024),

or by leverage. Similarly, a 25 basis point easing in the instrumented one-year Treasury yield raises automation adoption among constrained firms by 2 percentage points more than among unconstrained firms. This response is concentrated in the 2007–2013 period of falling rates and vanishes during the post-2013 normalization, consistent with the effect operating through the relaxation of financing conditions rather than rate movements per se.

A counterfactual exercise that holds the yen at its 2011 value implies that robot adoption among constrained firms would have been roughly 17% lower over 2012-2015, so a sizeable share of automation diffusion among financially constrained firms is attributable to the yen-driven decline in robot prices.

Entering both shocks jointly lets us weigh the two margins against each other, because the model below routes both through the same per-unit cash requirement. We map the elasticities from the empirical results to a financial primitive in the model, the advance rate, allowing for a quantitative check of the model against independent data. First, across firms, financial constraints amplify the price-shock response 56% more than the rate-shock response. Fed through the model, this implies that for the same piece of equipment a constrained firm must cover roughly 10 percentage points more of the hardware cost from internal funds (i.e. larger down payment). Second, across shocks, standardized by each shock's own identifying variation, the two differentials are of comparable magnitude, which implies an average advance rate of about 0.50 on level. These two objects align with firms' own equipment financing, moments the estimation never targets.

Event-study estimates around the yen depreciation show that the adoption gap between constrained and unconstrained firms opens within a year of the shock and persists for at least four years; instrumental-variable local projections deliver the same persistence pattern for the rate shock. Placebo tests using general capital expenditure and R&D, investment categories not directly exposed to imported robot prices, show no differential pre-trends and only weak, imprecise post-shock responses, in sharp contrast to automation. Since any generic trend in constrained firms' investment should be reflected in these broad categories rather than on the narrow automation-import margin, the flat placebo response rejects a confounding-trend or balance-sheet interpretation and favors the input-price channel.

These results are robust to a number of checks. The constraint-amplification pattern itself holds across four independent measures of financial tightness – leverage, cash, and the text-based constraints of **?** and Linn and Weagley (2024) – and under continuous rather than binary definitions, so it does not hinge on any single proxy or threshold. It also holds across multiple measures of adoption intensity, including shipment counts and imports scaled by assets.

We also rule out alternative explanations. The pattern is not explained by differential investment opportunities (it survives controlling for intangible-adjusted Tobin's Q), by globalization strategy (controlling for firm-level outsourcing), by the largest firms (dropping industry "star" firms), or by the robot-intensive automotive sector (excluding NAICS 336). Nor do they reflect selection into adoption: the effect survives restricting to plausible adopters via an ex-ante propensity score and under inverse-propensity-weighted and matched samples, ruling out that the result is driven by delayed adoption among firms inherently more prone to automate.

In addition, we show that the price effect is specific to the yen among exchange rates — controlling for interactions with the broad dollar, the renminbi, and an emerging-market currency index leaves it unchanged — and it appears where ex-ante exposure predicts: firms with pre-existing Japanese supplier relationships and firms in industries where automation is more applicable respond more strongly. The rate effect is not instrument-specific, recurring with credit spreads and the excess bond premium. Inference is robust to the small number of aggregate shocks under Driscoll-Kraay errors, the wild cluster bootstrap, and Anderson-Rubin tests.

Our findings should not be confused with the ordinary price elasticity of investment, where robots get cheaper so firms buy more. In a standard investment model, financing frictions dampen investment, so a fall in the price of capital goods moves constrained firms less than unconstrained ones. We emphasize the reversal of this heterogeneity. In the same vein, it not a simple exchange-rate pass-through story since the yen depreciation lowered the dollar price for every U.S. importer alike. The mechanism instead lies in how new technologies, particularly robots are financed. The hardware—arms, controllers, sensors—is standardized, redeployable, and traded in active secondary markets, making it pledgeable collateral in the sense of Benmelech et al. (2022). But turning hardware into

productive capacity requires complementary spending on software, integration, training, and process redesign—outlays lenders will not advance against, which can account for as much as three-quarters of a robot's installed cost (Leigh and Kraft, 2018).[2] Because this intangible component must be paid from internal funds, internal liquidity and not the sticker price, is the binding margin for many adopters. A fall in the hardware price relaxes that margin, and more so for firms whose cash was the binding constraint.

We formalize this financing structure in a collateral model featuring risky debt and non-pledgeable intangible costs. Firms finance automation with internal cash plus collateralized borrowing, when these two sources do not allow the firm to reach its unconstrained investment scale, the firm is finance-binding and adoption is pinned by the cash identity. In that regime, every friction collapses into a single per-unit cash requirement (or invoice, M) – the sticker price net of what can be borrowed against it, the intangible setup cost, and an installation term that rises with scale – and adoption is pinned by available cash divided by this invoice. Both elasticities follow in closed form: the price elasticity scales with the cash share of the per-unit cost (the exposure to sticker-price movements), and the rate elasticity with the financed share, which contracts as rates rise and the lender's advance shrinks.

The model delivers a falsifiable cross-equation restriction: the firms most responsive to capital-goods prices must also be the most responsive to financing costs, with both responses rising in balance-sheet tightness. The natural competing account is unobserved cost heterogeneity—some firms simply face lower integration and implementation outlays—and elasticities alone cannot make the separation: cheaper implementation, like tighter liquidity, raises both responses. The model instead delivers a sharper identification, in the spirit of sufficient-statistic approaches (Chetty, 2009) and the transparency arguments of Andrews et al. (2020b,a). First, low liquidity raises both elasticities *and* the shadow value of funds, while low implementation costs raise the elasticities but *relax* financial stress—so finding the largest responses precisely among firms independently classified as stressed

[2]Engineering sources likewise emphasize that installation, programming, and line re-engineering materially increase total expenditures beyond the hardware price (Engineering.com, 2018). 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. (2025), Liu et al. (2022)).

favors liquidity over a pure low-cost explanation. Second, every per-unit cost component cancels from the ratio of the two elasticities, which depends only on the lender's advance against the hardware; unobserved cost heterogeneity therefore cannot generate the cross-shock differential pattern we estimate. We take both diagnostics to the data.

This paper makes three contributions. First, on the empirical side, we document a reversal of the standard financial-frictions prediction: constrained firms respond more strongly than unconstrained firms to both favorable shocks, and—within the same firms and the same asset—the constrained–unconstrained differential is 56% larger for capital-goods prices than for financing costs. Establishing this ranking requires separating two forces most designs cannot, and it recasts financing frictions as something that can amplify, not merely dampen, responses to favorable shocks. On the financing-cost side this places us in a contested literature—Ottonello and Winberry (2020) find that low-risk firms respond more to monetary policy, while liquidity- and age-based sorts find the opposite (Cloyne et al., 2023; Jeenas, 2026); our design attributes the disagreement to the asset: when the capital good is pledgeable and front-loads cash, the constrained respond more.

Second, on the theoretical side, we develop a collateral model that links the two responses through a common liquidity margin and yields closed-form elasticities. Its central restriction is that exposure to price shocks and exposure to financing-cost shocks must move together across firms. That restriction is falsifiable, and ours is one of the few settings where it can actually be tested, because the two components of the user cost move separately. Because the model is closed-form in a single budget identity, it maps the reduced-form responses into financing primitives without estimating production parameters. The model provides a diagnostic that separates a liquidity mechanism from one driven by low intangible setup costs. Once we map the estimated elasticities through the model, they recover the advance rate lenders extend against robot hardware, at a level near one-half for constrained firms with a cross-firm gap of roughly ten percentage points—both matching firms' observed equipment financing, moments never targeted in estimation.

Third, on the measurement side, we build a transaction-level measure of automation adoption from bill-of-lading import records linked to Compustat, observing the timing, type, and origin of each

purchase. This resolves a recurring measurement problem: adoption has been inferred from capital-expenditure aggregates or industry-level robot counts, neither of which identifies which firm adopts what, and when. To our knowledge, this is the first transaction-level, firm-level panel of automation adoption for the United States linked to financial statements; it complements industry-level IFR penetration measures, survey-based adoption waves, and customs-based firm studies for European economies (Koch et al., 2021; Bonfiglioli et al., 2024).

More broadly, our findings imply that macroeconomic shocks and trade-policy movements affect not only the aggregate pace of automation adoption, but also which firms automate. By disproportionately relaxing adoption constraints for financially constrained firms, favorable price shocks can compress technological gaps across firms, while tighter financing conditions can slow diffusion among firms operating near liquidity limits. This is a collateral-and-liquidity channel through which macroeconomic shocks govern the cross-sectional diffusion of new technologies.

Most work on industrial robots focuses on what happens after adoption: labor displacement (Acemoglu and Restrepo, 2020; Dauth et al., 2021; Bessen et al., 2025), productivity gains (Graetz and Michaels, 2018; Koch et al., 2021; Rodrigo, 2023), and the concentration of robots in a small number of industries and regions (Benmelech and Zator, 2025; Brynjolfsson et al., 2023). Yet aggregate robot investment in the U.S. remains small—under 0.3% of equipment spending—and highly uneven across firms, raising a prior question that has received far less attention: what determines which firms adopt in the first place? As argued above, an important part of that answer lies in the financing structure of the capital good, alongside the labor-market and productivity motives this paper emphasizes: because robots bundle pledgeable hardware with non-pledgeable intangibles, the equipment's price and financing cost sort firms along balance-sheet lines.

Our findings also speak to a growing body of work on heterogeneous responses to monetary policy. Kroen et al. (2025) show that falling interest rates disproportionately benefit industry leaders whose value derives from long-duration intangible assets—brands, patents, organizational capital. In that setting, the marginal beneficiaries of rate cuts are large, unconstrained firms at the top of the size distribution. We identify the opposite regime: when the relevant capital is pledgeable and

front-loads cash outlays, constrained followers, not unconstrained leaders, are the marginal adopters, and the dominant margin is the capital-goods price—which moves with exchange rates and trade policy—rather than the policy rate. The two channels, duration for intangibles and collateral for pledgeable equipment, operate on different asset classes and different parts of the firm distribution, and can naturally coexist. Which channel dominates in a given context depends on the financing structure of the capital good in question, a distinction our framework makes precise.

Although we study industrial robots, the financing structure we identify, i.e., pledgeable equipment bundled with non-pledgeable intangibles, characterizes a broad class of capital intensive technologies (e.g., 3D printers, MRI and CT scanners requiring facility retrofits and technician training, enterprise software implementations, etc.). Our findings therefore speak to technology diffusion more broadly: for capital goods featuring this financing structure, capital-goods-price shocks may dominate financing-cost shocks in driving adoption among constrained firms, with implications for how trade policy[3] and exchange-rate movements shape the pace and distribution of technological upgrading.

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 estimates the causal effects of yen and rate shocks. Section 6 maps the results back to the model and calibrates the advance rate from the estimated elasticities, and Section 7 concludes.

2 Literature

Our work connects four strands of research: financing frictions and investment, collateral and pledgeability, new technologies and intangible investment, and heterogeneous transmission of price and macro shocks to firm investment.[4] The classic view is that financial frictions prevent constrained firms from scaling into profitable opportunities.[5] With borrowing frictions or scarce collateral,

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

[4] Internet Appendix IA.1 provides an overview of related work organized by measurement, shock variation, and mechanism.

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

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). Our results point in the opposite direction: constrained firms respond *more* to both price and rate shocks. The difference arises because robot capital is partially pledgeable—when its price falls or borrowing costs decline, the binding cash constraint relaxes most for firms at the collateral margin.[6]

Our model builds on Rampini and Viswanathan (2010, 2013), who show that constrained firms optimally conserve debt capacity rather than fully collateralizing assets. We adapt their logic by collapsing all robot-related expenditures into a per-unit invoice that splits the cost into a pledgeable purchase price and a non-pledgeable intangible component. This decomposition makes transparent how each cost component maps into collateral capacity and therefore into the user cost relevant for marginal adopters. The hybrid-style costs are a key ingredient: the response to the shock is the same across constrained and unconstrained firms under fully pledgeable costs.

Our analysis is complementary to Almeida et al. (2004); Almeida and Campello (2007) and Bolton et al. (2011), who study the roles of liquidity management, precautionary savings, and collateral multipliers in shaping dynamic investment and financing choices. Rather than studying broad investment capacity, we focus on per-unit technology adoption decisions and show how user-cost shocks load differently onto the pledgeable and non-pledgeable components of a robot's cost—generating a mapping from cost composition to heterogeneous adoption elasticities that we take directly to the cross-sectional evidence.

Recent work has refined our understanding of how collateral constraints bind. Lian and Ma (2021) document that 80% of U.S. corporate lending is cash-flow-based rather than asset-based, although asset-based constraints become more binding during downturns. For automation equipment, however,

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

asset-based lending is the natural benchmark: robot hardware is standardized, redeployable, and traded in active secondary markets, precisely the characteristics that sustain borrowing-base facilities (Benmelech et al., 2022). The complication is that robots are not just hardware. Converting a robot arm into productive capacity requires integration, training, software customization, and process redesign—costs that can account for three-quarters of total installed system costs (Leigh and Kraft, 2018) and that no lender will advance against. Our model captures this dual structure: the pledgeable hardware can be financed through asset-based borrowing, but the non-pledgeable intangible component must be covered out of internal cash, creating the cash-envelope friction that drives our results. Chaney et al. (2012) and Catherine et al. (2022) show that collateral revaluations affect firm investment through a borrowing channel, and Drechsel (2023) provides complementary evidence that firms facing tighter earnings-based constraints respond more strongly to monetary policy. In both cases the mechanism operates through the overall tightness of the borrowing environment. Our contribution is to decompose the cost of a single investment into its pledgeable and non-pledgeable parts and show that this composition, rather than the level of collateral values or the severity of borrowing constraints, determines which firms respond most to price and rate shocks. The rising importance of intangible capital documented by Falato et al. (2022) suggests that this structure extends well beyond robots.

Our model also relates to work showing that falling rates can favor large, intangible-intensive "superstars" via duration and easy finance (Kroen et al., 2025; Liu et al., 2022). 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, 2026). 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.

Most empirical work on robots studies consequences of adoption—labor displacement (Acemoglu and Restrepo, 2020; Dauth et al., 2021; Bessen et al., 2025), productivity (Graetz and Michaels, 2018; Koch et al., 2021; Rodrigo, 2023), and geographic concentration (Benmelech and Zator, 2025; Brynjolfsson et al., 2023). The determinants question has received less attention, despite the fact that aggregate robot investment remains small and highly uneven across firms. Our research addresses this gap by showing how macroeconomic conditions, specifically exchange rates and interest rates, can reallocate adoption toward financially constrained firms, thereby broadening access to productivity-enhancing technologies.[7] In addition, our findings align with micro evidence from the 2019 American Business Survey analyzed by Acemoglu et al. (2022), in which firms most frequently cite that robots are "too expensive" and that they face a "lack of capital" as reasons for not adopting.

We measure adoption from transaction-level bills of lading linked to Compustat, which delivers firm-level timing on the adoption margin. Comparable firm-level approaches in European data include Bonfiglioli et al. (2024); Acemoglu et al. (2020); Humlum (2022); U.S. survey and IFR-based evidence appears in Bessen et al. (2025); Koch et al. (2021); Antonioli et al. (2024); Brynjolfsson et al. (2023). Our identification leverages exchange-rate pass-through to imported input prices (e.g., Amiti et al., 2014; Blaum et al., 2018; Gopinath and Neiman, 2014; Di Giovanni et al., 2022) and high-frequency monetary policy surprises following Döttling and Ratnovski (2023), using the orthogonalized series of Bauer and Swanson (2023) to mitigate Fed-information confounds (Gertler and Karadi, 2015; Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021).

3 Model

Our central empirical finding is a puzzle for the standard view of financing frictions. When the yen depreciation and the monetary-policy surprises lowered the user cost of automation, the firms that responded most were the constrained ones — those with the least internal cash, the weakest collateral capacity, and the most binding debt limits. A useful model must therefore do three things. First, explain why financing constraints amplify rather than dampen adoption. Second, deliver a falsifiable joint restriction — a sign pattern linking the price and rate responses across the cross-section —

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

rather than a single elasticity. Third, map directly onto a regression we can estimate on transaction-level data without recovering production parameters. We develop a two-period collateral model that meets all three: it abstracts from production estimation by focusing on financing feasibility when constraints bind, expresses responses to price and rate shocks through a transparent sources–uses identity, and aligns with the short-run exogenous shocks we exploit (exchange-rate pass-through to robot prices, p_k and high-frequency monetary surprises to interest rates, R).

3.1 Economic Environment

Time is discrete with two periods, $t \in \{0, 1\}$. At $t = 0$ a firm with internal cash c chooses robot capital k', external borrowing b, and an encumbrance share $\theta \in [0, \bar{\theta}]$. All decisions are made before uncertainty resolves.

Robot investment combines two components. The first is tangible hardware purchased at market price p_k. The second is a bundle of complementary intangible expenditures, h—software, integration, training, and production-line reconfiguration—required to make the hardware productive. Industry evidence indicates these integration costs are at least as large as hardware costs and often substantially larger. We treat h as a per-unit non-pledgeable component, capturing the dual structure of robot investment: pledgeable hardware bundled with non-pledgeable intangibles. Scaling adoption also entails convex installation cost $\frac{\tau}{2}(k')^2$, reflecting adjustment frictions.

At $t = 1$, uncertainty resolves. With probability π, the firm survives, produces $A(k')^\alpha$, retains undepreciated asset value $(1 - \delta)p_k k'$, and repays b' to lenders. Equity's payoff is

$$\Pi_1^{\text{good}} = A(k')^\alpha + (1 - \delta)p_k k' - b'. \tag{3.1}$$

With probability $1 - \pi$, the firm defaults and assets liquidate for $\rho p_k k'$, with $\rho \in (0, 1)$. Creditors hold a legal claim on fraction θ of assets but recover only $m(\theta) \leq \theta$ due to enforcement frictions, with the remainder dissipated as deadweight loss. Equity receives

$$\Pi_1^{\text{bad}} = (1 - \theta)\rho p_k k' \tag{3.2}$$

while creditors receive $m(\theta)\rho p_k k'$.

The encumbrance choice θ 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 Financing Structure

Lenders are risk-neutral and discount at gross policy rate R. Under competitive pricing, the zero-profit condition implies:

$$\frac{\pi b' + (1-\pi)m(\theta)\rho p_k k'}{R} = b \tag{3.3}$$

where b' is the promised repayment. This defines the contractual rate $r_c \equiv b'/b = r_c(R)$ with $r_c'(R) > 0$ and $r_c'(R) < 1/\pi$, ensuring monetary tightening reduces borrowing capacity.

Borrowing capacity follows a reduced-form constraint:

$$b \leq y(R)\, m(\theta)\, p_k k' \tag{3.4}$$

where $y(R) \equiv (1-\pi)\rho/(R - \pi r_c)$ is the advance rate per dollar of effective pledgeable collateral. Monetary tightening reduces advances: $y_R < 0$.

The sources-uses budget identity is:

$$c + b = (p_k + h)k' + \tfrac{\tau}{2}(k')^2 \tag{3.5}$$

A feasible choice (k', b, θ) satisfies equations (3.5), (3.4), and the cap constraint, with $k' \geq 0$, $b \geq 0$, and $0 \leq \theta \leq \bar{\theta}$.

3.3 Role of Cash

For cash-rich firms where financing is slack, investment is determined by standard marginal-product conditions and cash does not affect scale (see Appendix IA.4.1 for derivation of the elasticities under this regime).

Thus, the differential responses we estimate are identified by cash-poor firms in the binding regime. Because scale is pinned by the cash identity rather than production-function curvature, we can recover elasticities without estimating production parameters.

3.3.1 Cash Identity for Constrained Firms

For cash-poor firms where both budget and borrowing constraints bind, substituting the binding borrowing condition $b = y\, m(\theta)\, p_k k'$ into the budget identity (3.5) yields the fundamental cash identity:

$$k' \cdot M = c \tag{3.6}$$

where the per-unit cash requirement M is:

$$M \equiv p_k(1 - y\, m(\theta)) + h + \tfrac{\tau}{2}k' \tag{3.7}$$

This identity compresses all financing frictions into a single reduced form: each unit of capital requires internal cash M, so scale is pinned by $k' = c/M$ whenever financing binds. Figure 1 decomposes M into three components: hardware cost net of financing $p_k(1-y\, m(\theta))$ (the "exposure"), non-pledgeable intangibles h, and installation costs $\frac{\tau}{2}k'$ (the "brake"). See the step by step derivation of the cash identity in Appendix IA.4.2.

Because scale is governed by the invoice M rather than the curvature of the production function, the elasticities are estimable without recovering firm-level productivity, and the coefficient on each shock interaction identifies the corresponding elasticity directly. This permits a joint test of the price and rate channels on a common adoption margin that aggregate capital-expenditure data cannot support.

3.4 Investment Elasticities

Log-differentiating the cash identity $k'M = c$ with respect to price and rate shocks yields closed-form elasticities. Define exposure shares: $W_{p_k} \equiv p_k(1 - y\, m(\theta))/M$ (hardware net of financing), $W_f \equiv p_k\, y\, m(\theta)/M$ (financed share), and the installation share $\kappa \equiv (\tau/2)k'/M$.

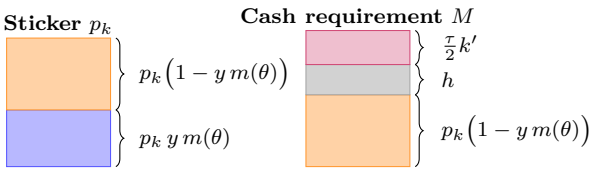


Figure 1: Decomposition of hardware sticker price p_k and the per-unit cash requirement M. The blue portion of hardware can be financed; the orange portion must be paid in cash. Together with non-pledgeable intangibles h and installation costs, this defines the total cash requirement per unit.

There are two subcases of cash-poor firms to consider. Cash-poor firms that are "at the cap" are those whose encumbrance 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.

Proposition 1 (Investment Elasticities at the Collateral Cap)**.** *For cash-poor firms at the collateral cap ($\theta^* = \bar{\theta}$), the price and rate elasticities are:*

$$\varepsilon_{k',p_k} = -\frac{W_{p_k}}{1+\kappa} < 0, \qquad \varepsilon_{k',R} = \frac{W_f\, \chi_R}{1+\kappa} < 0 \tag{3.8}$$

where $\chi_R \equiv \partial \ln y / \partial \ln R < 0$ captures lending pass-through.

Proof in Appendix IA.4.3. The intuition is straightforward. When hardware prices fall, the cash requirement M declines proportionally to the exposure share W_{p_k}—the fraction of costs that must be paid in cash. When borrowing costs rise, the financed share contracts by $W_f \chi_R$ as lender advances decline. Comparing the two, the price elasticity dominates whenever the out-of-pocket share W_{p_k} exceeds the rate-weighted financed share $W_f|\chi_R|$. Both effects are dampened by installation frictions (the brake term $1+\kappa$). Critically, firms with tighter balance sheets—less internal cash, weaker collateral capacity—have higher W_{p_k} and W_f, producing larger responses to both shocks. For unconstrained firms the shock operates on the intensive margin of profitability; for constrained firms it relaxes the binding budget directly, so the same internal cash finances a larger robot stock.

Proposition 2 (Investment Elasticities with Interior Pledge)**.** *For cash-poor firms with interior pledge ($0 < \theta^* < \bar{\theta}$), the elasticities are:*

$$\varepsilon^{int}_{k',p_k} = -\frac{W_{p_k} - W_\theta\, \eta_{\theta,p}}{1+\kappa} < 0, \qquad \varepsilon^{int}_{k',R} = \frac{W_f\, \chi_R + W_\theta\, \eta_{\theta,R}}{1+\kappa} < 0 \tag{3.9}$$

where W_θ is the pledgeability margin and $\eta_{\theta,x}$ captures how encumbrance responds to shock x.

Proof in Appendix IA.4.4.

Pledge adjustment dampens price sensitivity (as firms can borrow more when hardware is cheaper) but can amplify or mitigate rate sensitivity depending on how encumbrance responds. In both regimes, however, the key insight remains: both elasticities load on the same balance-sheet components, generating testable joint restrictions.

3.5 Distinguishing Liquidity from Intangible Costs

A critical tension emerges when interpreting cross-sectional variation in elasticities. Robots differ not only in how constrained firms are (internal cash c) but also in the complementary intangibles required for integration (customization costs h). Both primitives affect the cash requirement M in ways that could produce similar elasticity patterns, yet they reflect fundamentally different economic mechanisms.

Lower internal cash ($c \downarrow$) tightens the financing constraint: the firm has less liquidity to cover the non-pledgeable components of adoption. Lower intangible costs ($h \downarrow$) reduce the cash burden per unit, making adoption cheaper in a way that mechanically resembles having more cash available. From equation (3.7), a fall in either c or h can raise the exposure shares W_{p_k} and W_f, potentially increasing both elasticity magnitudes $|\varepsilon_{k',p_k}|$ and $|\varepsilon_{k',R}|$. Elasticities alone therefore cannot distinguish tight liquidity from low setup costs.

The resolution lies in the joint behavior of elasticities and financial stress indicators. At the collateral cap, changes in primitives affect the shadow value of funds λ (the Lagrange multiplier on the budget constraint). Inside the cap, they affect the optimal pledge θ^*. Table 1 summarizes the key predictions.

The first two rows isolate the key distinction. A fall in cash c raises both elasticities *and* increases financial stress (higher λ at the cap, higher θ^* in the interior), as the firm becomes more liquidity-constrained and must encumber more assets to finance adoption. A fall in intangibles h also raises both elasticities but *decreases* financial stress (lower λ, lower θ^*), as the reduced cash burden per unit relaxes the financing constraint and reduces the need for pledging.

Table 1: Comparative statics: Distinguishing liquidity from intangible costs

Primitive change	$\lvert\varepsilon_{k',p_k}\rvert$	$\lvert\varepsilon_{k',R}\rvert$	Financial stress (λ or θ^*)
Cash $c \downarrow$ (tighter liquidity)	↑	↑	↑
Intangibles $h \downarrow$ (cheaper setup)	↑	↑	↓
Install frictions $\tau \downarrow$	↑	↑	↓
Collateral cap $\bar{\theta} \downarrow$	↑	↓	↑

Notes. At the collateral cap, financial stress is measured by λ (shadow value of funds). In the interior regime, it is measured by θ^* (optimal encumbrance). Both $c \downarrow$ and $h \downarrow$ raise elasticity magnitudes, but only $c \downarrow$ increases financial stress. Lower intangibles relieve the cash burden, reducing stress. Full comparative statics in Appendix IA.4.5, Tables IA.4.1–IA.4.2.

This produces a falsifiable joint restriction. Under the model assumptions, observing that both elasticities are larger precisely among firms that exhibit higher financial stress—measured through low cash-to-assets ratios, high debt-constraint language in 10-K filings (**?**Linn and Weagley, 2024), or high leverage—points to internal liquidity c rather than low intangible costs h as the primary driver of cross-sectional heterogeneity.[8]

The lower rows of Table 1 show that other primitives produce distinct signatures. Lower installation frictions ($\tau \downarrow$) resemble cheaper intangibles: both elasticities rise but stress falls. A tighter collateral cap ($\bar{\theta} \downarrow$) raises price sensitivity but *lowers* rate sensitivity, as the firm cannot borrow more even when rates fall—a pattern inconsistent with what we observe empirically. The diagnostic therefore allows us to rule out alternative explanations based on heterogeneity in adjustment costs, collateral capacity, or lending terms, all of which imply different co-movement patterns than what the data deliver.

The comparative statics above concern the *levels* of the two elasticities, which move with any primitive entering the invoice. Their *ratio* does not, allowing for a separate strategy to check the consistency of the model with observed financing patterns.

Corollary 1 (Cost invariance of the cross-shock ratio)**.** *At the collateral cap, the elasticities of Proposition 1 satisfy*

$$\left|\frac{\varepsilon_{k',p_k}}{\varepsilon_{k',R}}\right| = \frac{W_{p_k}}{W_f\,|\chi_R|} = \frac{1-\varphi}{\varphi\,|\chi_R|}, \qquad \varphi \equiv y\,m(\bar{\theta}), \tag{3.10}$$

[8] Formally, the triple $\left(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \lambda\right)$ at the cap or $\left(|\varepsilon_{k',p_k}|, |\varepsilon_{k',R}|, \theta^*\right)$ in the interior uniquely identifies a liquidity tightening ($c \downarrow$) among the class of single-primitive shocks in our model. See Appendix IA.4.8 for the formal statement and proof.

independent of the intangible cost h, the installation term $(\tau/2)k'$, internal cash c, and the level of the hardware price p_k.

The result is immediate from Proposition 1: both elasticities carry the common transmission factor $1/(1+\kappa)$, and the invoice M cancels from the ratio of the loadings, $W_{p_k}/W_f = [p_k(1 - y\,m(\bar{\theta}))/M]/[p_k\,y\,m(\bar{\theta})/M] = (1-\varphi)/\varphi$. [Check before circulation: the analogous statement in the interior regime, with $\varphi^* \equiv y\,m(\theta^*)$, should follow from Proposition 2 by the same cancellation; verify against the interior expressions before asserting it in text.]

Here φ is the *advance rate*, or loan-to-value: the fraction of the hardware price a lender will finance against the machine as collateral, with the firm paying the rest from its own funds. The corollary says this single financing parameter governs how automation responds to a cheaper machine *relative* to cheaper credit—a lower price relieves the cash-paid share $1-\varphi$, a lower rate relieves the financed share φ, so a firm that funds more of the machine out of pocket is more price- than rate-sensitive. Everything else about the purchase—the integration cost h, installation, scale, and internal cash—moves *both* elasticities together and cancels from the ratio, leaving only the financing margin. This ratio can be identified in our empirical exercise, creating a useful check on the model. A firm's price-versus-rate sensitivity reveals the advance rate it faces, which we can also measure directly from how firms fund their capital purchases. Section 6.3 carries out this calibration for firms on the finance-binding margin.

3.6 Empirical Predictions

The model generates three testable predictions. First, constrained firms, i.e., those with lower internal cash, weaker collateral capacity, or higher non-pledgeable costs, should exhibit larger adoption elasticities to both price and rate shocks, with $|\varepsilon_{k',p_k}|$ and $|\varepsilon_{k',R}|$ both increasing in financial tightness. Second, the relative magnitude of the two responses is governed by the financing split rather than by any cost primitive: by Corollary 1, $|\varepsilon_{k',p_k}/\varepsilon_{k',R}| = (1-\varphi)/(\varphi\,|\chi_R|)$, so the price loading dominates exactly when the out-of-pocket share exceeds the financed share scaled by lending pass-through—an object Section 6.3 recovers from the data rather than assumes. Third, the joint restriction—firms that respond strongly to one shock should also respond strongly to the other—rules out explanations based

on collateral capacity, which imply opposite-signed co-movement; the pairing with financial-stress indicators and the cost-invariant ratio then separate the liquidity mechanism from heterogeneous intangible costs.

The comparative statics also provide a diagnostic for the underlying friction. Lower internal cash c raises both elasticity magnitudes and increases the shadow value of funds λ (at the cap) or optimal pledge θ^* (in the interior). In contrast, lower intangible costs h also raise both elasticities but *decrease* λ and θ^*. The challenge is that elasticities alone cannot distinguish these two drivers—both tighten the effective budget constraint. However, Andrews et al. (2020b) argue that reporting informative auxiliary moments or descriptive statistics can help readers recover the economic mechanisms underlying structural estimates, a point further developed in Andrews et al. (2020a). Our diagnostic exploits this insight: 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. Observing larger elasticities precisely among firms that appear more financially stressed points to internal liquidity rather than low setup costs as the underlying driver. Formal comparative statics and diagnostic tests are provided in Appendix IA.4.5 and IA.4.6.

These predictions map directly to our empirical strategy. We measure financial constraints through multiple ex-ante proxies: cash holdings, leverage, and debt-covenant language. For the diagnostic test, we do not observe λ directly, but we use text-based measures of firms' disclosures about liquidity constraints as proxies. Specifically, we 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.[9] Conceptually, such statements correspond to a high λ in our model: the marginal value of relaxing the financing constraint is large because otherwise profitable projects must be postponed. We then exploit independent variation in robot prices (yen depreciation) and financing costs (monetary policy

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

surprises) and test whether firms with high debt-constraint scores - indicating high λ - respond to both shocks with larger elasticities, as predicted. Section 5 translates these predictions into the estimating equation, interacting the price and rate shocks with predetermined measures of financial tightness, and reports the associated diagnostic tests.

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,[10] 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 (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).[11]

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.2.1. From the raw import

[10] 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).

[11] 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)).

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.2.3 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.[12] 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. Although robot purchases are recorded at a high frequency, automation projects are scoped, budgeted, and approved on an annual capital-budget cycle, while procurement takes several months. We thus collapse robot transactions to the annual level, which aligns measurement with the actual investment decision process.

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 surveys of industrial robot suppliers worldwide,[13] 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.

[12]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.

[13]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."

Since IFR data exclusively covers industrial robots rather than broader automation technologies, Figure 2 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. The higher correlation for the broader measure reflects the fact that aggregating across related HS codes smooths shipment-level noise and aligns more closely with IFR's system-level reporting. These strong and robust correlations affirm the validity of our novel measurement approach.[14]

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 2 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).[15] To construct our final sample, we first aggregate the shipments to the parent 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 Compustat sample, with 5,442 firm-year

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

[15] BECRS provides crosswalks across different firm identifiers including links between Capital IQ identifier (*companyid*) and Compustat identifier (*GVKEY*). See Table IA.2.1 for matching success.

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 4 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.[16]

FIGURE 4 ABOUT HERE

[16] 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.

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.[17] Panel A of Table 2 lists the countries from which most industrial robots shipped to the U.S. originate, while Panel B focuses on automation technologies more broadly. Both panels reveal strong concentration of imports from a small set of countries, with five countries accounting for approximately 81% of the volume and 72% of the shipments of imported robots. Japan emerges as the dominant supplier, accounting for 53% of total robot import volume, followed by Germany at 14%. This Japanese dominance reflects the country's historical leadership in robotics technology and manufacturing.

TABLE 2 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 2 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.[18]

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 3 shows the sectoral distribution of automation and robot adoption among publicly traded firms in our sample.

[17] Our country of origin is produced in the same way as described in Barbiero et al. (2026), where the Panjiva field *placeofreceipt* indicates the location where the shipment originates. We use a fuzzy matching algorithm to pick out country names from this field. In the cases where this field is empty, we default to *shpmtorigin*, which indicates the country from which the shipment departed for the United States.

[18] 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.

The data reveal that automation adoption is heavily concentrated in specific sectors. Manufacturing emerges as the dominant sector for automation adoption. Measured as an annual adoption probability, 24% of manufacturing firm-year observations had automation technology shipments and 7% had industrial robot shipments. 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 are in manufacturing. 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. At this level, we report cumulative "ever-adopter" shares, defined as the fraction of firms within a subsector that import at least one automation technology over the sample period. The data reveal substantial heterogeneity across manufacturing subsectors. Metal Products and Machinery emerges as a leading subsector, with 48% of firms ever adopting automation technologies and 8% ever adopting industrial robots. This subsector accounts for 17.44% of all automation adopters and 20.99% of robot adopters within manufacturing. The high adoption rates likely reflect production processes involving repetitive metalworking operations well-suited to robotic automation, such as welding, cutting, and assembly.

Computers and electronics shows similarly high within-subsector adoption for automation technologies (25%) but lower robot adoption rates (4%). Despite the lower robot adoption rate, this subsector represents a substantial 20.62% of automation adopters and 20.8% of robot adopters within manufacturing, reflecting the scale of the sector and the importance of precision and clean-room automation in electronics production. The Automotive subsector exhibits the highest robot adoption intensity, with 48% of firms ever adopting automation technologies and 14% adopting industrial robots—the highest robot adoption rate among manufacturing subsectors. Although Automotive accounts for only 8.54% of automation adopters, it represents a disproportionately large 17.41% of robot adopters within manufacturing. This pattern reflects the industry's pioneering role in industrial robotics, particularly in welding, painting, and assembly operations.

Panels B and C of Table 3 list the top ten U.S. companies adopting industrial robots and automation, 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 3 ABOUT HERE

4.3 Automation Measures, Financial Constraints, and Policy Shocks

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 automation-related shipment entries, denoted *Shipments*, where each shipment entry corresponds to a single automation-related HS6 product line recorded in Panjiva and is derived from Bills of Lading, which may list multiple products within a single document. 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.[19]

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

[19]Each Bill of Lading (BoL) may list multiple products at the HS6 level. We classify a BoL as automation-related if any listed HS6 code corresponds to an automation good. Multi-product BoLs - those listing more than one HS6 code - account for only about 10 percent of imports and typically involve closely related HS categories. For volume-based measures, we allocate container volume (TEUs) proportionally across products listed on the BoL. Approximately 5 percent of shipment records lack TEU information, usually because the shipment is not containerized. Over the sample period, the correlation between *Shipments* and *Volume* is 0.98.

assets *(ROA)*, calculated as net income divided by total assets.

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. Financing constraints are measured with 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.[20] For periods for which both measures are available, 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. As robustness measures, *High FC (HM)* equals one based on the Hoberg-Maksimovic constraint measure and *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.

Exchange Rate Shocks As reported in Table 2, Japan, the dominant global producer of industrial robots,[21] 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 monetary easing to combat deflation (commonly referred to as "Abenomics")—marked a significant

[20] 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.

[21] According to the International Federation of Robotics, Japanese manufacturers delivered 45% of the global supply of industrial robots in 2022.

break from previous policy, triggering exchange rate movements largely unrelated to U.S. macroeconomic conditions or firm-level factors. Figure 3 plots the Yen–USD exchange rate. Between 2012 and 2015, the Yen depreciated approximately 33%, substantially reducing the dollar cost of Japanese robots. Over the same period, the BoJ's Export Price Index for 'General purpose, production and business oriented machinery' rose 17% on a yen basis.[22] Thus, the rise in yen-denominated prices is far smaller than the yen depreciation, implying a substantial decline in the dollar price of Japanese machinery exports. In Appendix Figure IA.2.1 we plot the U.S. import price index for Japanese Machinery (NAICS 333) along with the U.S. domestic Machinery PPI (as a benchmark for machinery prices). The import price index for Japanese goods is only available starting in 2012, but it drops over 10% from 2012-2015, with a clear deviation from the US Machinery PPI that displays a secular upward trend.

FIGURE 3 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. Because supply is concentrated and Japanese producers set the reference price, a yen-driven cost decline lowers the market price of automation for all U.S. buyers—partly through cheaper Japanese equipment and partly through competitors' price responses—so the relevant treatment is a common movement in the price of automation rather than a buyer-specific bilateral price; Section 5.2.2 develops and tests this distinction. 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.

To verify that our results are not driven by general dollar strength, we include the *BroadUSD* index

[22] Source: Quarterly data from the Bank of Japan, found here: BOJ; Series PRCG15_2400420001.

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.

Monetary Policy Shocks Automation investment entails 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).[23] 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

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

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 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 4. 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.[24] Furthermore, 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.

[24]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.

TABLE 4 ABOUT HERE

Table 4 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 robust evidence that firms with higher leverage ratios are more likely to adopt automation technologies, suggesting that leveraged firms may use automation as a strategy to improve operational efficiency and service their debt obligations more effectively.

Finally, the relationship between profitability and automation adoption is negative and significant across all measures: less profitable firms are more likely to adopt automation. We emphasize that these cross-sectional associations are descriptive and are not a test of the model, whose predictions concern the *elasticity* of adoption to price and financing shocks rather than adoption levels conditional on size, age, and leverage. If anything, the negative coefficient is consistent with the profile of the financially constrained "marginal adopter" that drives our shock analysis below: low current profitability is one observable marker of the low-internal-funds firms for whom automation is a margin that responds to relaxations in cost and financing constraints.

Overall, the results show that automation adoption is characterized by firm scale, maturity, and financial leverage 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 2, 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 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, $Shock_t$ denotes the relevant macroeconomic shock, and we include fixed effects for both firm and year to control for unobserved heterogeneity and aggregate trends. In the exchange-rate specifications, $Shock_t$ is the contemporaneous change in the yen–dollar exchange rate. In the interest-rate specifications, $Shock_t$ is the lagged fitted change in the one-year Treasury yield, instrumented using high-frequency monetary policy surprises. We also include the baseline firm controls (size, age, profitability, and leverage) in X_{it}. Standard errors are clustered at the firm level.[25]

[25] Our coefficient of interest captures differential firm responses to aggregate shocks based on predetermined constraint status. We cluster standard errors at the firm level to account for serial correlation in firm-level unobservables. To address potential correlation across firms induced by the common annual shock, Table IA.3.7 reports Driscoll–Kraay standard errors and wild-cluster-bootstrap p-values at the year level; both confirm our results.

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 53% of U.S. industrial robot imports), yen movements provide direct variation in robot prices that is plausibly exogenous to U.S. firm characteristics. To state more explicitly, the identifying assumption is that annual yen innovations—driven by Japanese monetary policy—are as good as randomly assigned with respect to the *differential* automation demand of constrained versus unconstrained U.S. firms, conditional on year fixed effects and firm controls. Thus, we test directly the predictions of Propositions 1 and 2 that pertain to the price elasticity.

Panel A of Table 5 presents results where we do not include year fixed effects – but instead include 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 0.0424 IHS units (roughly 4.2%, column 2)[26], 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 5 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

[26]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.

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 below the median cash-to-assets ratio in their first year of sample appearance, exhibit significantly stronger automation responses to Yen depreciation than their low-constraint 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.

The magnitude of these differential effects estimated in Table 5 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).[27]

To ensure these effects reflect Japan-specific price changes rather than broader dollar strength, we

[27]The decomposition uses the differential effect and the fact that approximately 40% of the regression sample is constrained firms (we use the median value of liquidity across all firms in our Compustat sample, so the share of constrained firms in the regression sample is not necessarily half). The "unconstrained effect", E_U, can be solved from $0.6E_U = 0.225 - 0.4(\beta + E_U)$. With the 0.225 average effect (Panel A) and the differential, $\beta = 0.161$ (Panel B), the total effect from a Yen depreciation is 0.16 for unconstrained firms and 0.322 for constrained firms (twice as responsive). The same calculation for debt constraints yields differential responses from 1.6 to 2.1 times larger for constrained firms.

control for the interaction with other exchange rates as well in Appendix Table IA.3.1. Specifically, we repeat the baseline specification while including interactions with the Chinese Yuan, an index for emerging economy exchange rates, and the broad USD index. This test yields reassuring results: 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.

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

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

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 within-year differential path of financially constrained firms relative to unconstrained firms in event-time ℓ years from the 2011 baseline (with $\ell = 0$ excluded as the reference period). The specification includes lagged firm controls ($X_{i,t-1}$: log age, log assets, leverage, and ROA), calendar-year fixed effects (γ_t) and firm fixed effects (α_i). Year fixed effects absorb aggregate time-series movements common to all firms, so the (β_l) coefficients are identified from within-year differences between constrained and unconstrained firms, normalized to 2011.

This specification allows us to examine both the timing and persistence of differential automation responses while providing a clean test for pre-existing differential trends. 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 the subsequent implementation of "Abenomics" generated the sharp yen depreciation that serves as our price shock.

Figure 5 presents the estimated event-study coefficients for the differential effects of automation

adoption for financially constrained firms based on their cash liquidity levels.[28] 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 5a 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. Figure 5b suggests the same pattern holds for the relative intensity of firm automation.

To further validate that our results do not reflect spurious correlation with other trends, Figure 6 presents equivalent event studies for general capital expenditure (CAPX) and research and development (R&D) spending by cash constraints. Two features distinguish these placebos from the automation results. First, neither series exhibits a differential pre-trend: through 2011 the constrained–unconstrained gap is flat and statistically indistinguishable from zero, so the two groups entered the shock on parallel general-investment paths. Second, although both series turn modestly positive after 2012, the post-shock responses are smaller, noisier, and largely insignificant year-by-year. This pattern disciplines the set of admissible confounders. The flat pre-2012 placebos rule out the most natural alternative—that constrained and unconstrained firms were already on diverging trajectories in broad investment before the shock. A generic balance-sheet or "financial-relief" channel, which relaxes overall investment capacity, would predict a differential response at least proportional to total CAPX, which dwarfs automation imports in dollar terms; instead the large, precisely estimated response appears on the narrow automation-import margin directly priced in yen, while the broad

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

aggregates responses are muted and imprecise.[29]

FIGURES 5 AND 6 ABOUT HERE

Timing and Option Value The persistence of the event-study response also addresses a dynamic concern that our two-period model abstracts from. If firms expected the depreciation to persist, waiting could have been optimal, and the estimated differential could in principle reflect the re-timing of predetermined investment rather than net new adoption. Two observations weigh against this reading. First, re-timing predicts a differential that spikes and then reverses as pulled-forward purchases are repaid in lower subsequent adoption; instead the constrained–unconstrained gap opens within a year of the shock and persists for at least four years, with no offsetting negative response at longer horizons. Second, the option to wait is worth most to firms that can transact at a time of their choosing. This is not the case for a cash-constrained firm whose adoption at the pre-shock price lay outside its budget set. Selection on waiting would concentrate the measured response among unconstrained firms. A full treatment of adoption timing under stochastic prices is beyond the static model, but our persistence result indicates that the timing margin does not drive the differential.

5.2.2 Where is the effect Concentrated? The Price Channel across Origins and Industries

A remaining question is whether the constraint–yen interaction identifies a Japan-specific robot-price shock or simply a broad capital-cheapening episode coinciding with 2012–2015. At the aggregate level, Appendix Figure IA.3.1 shows that Japan-origin industrial robot shipments in the full Panjiva import universe roughly doubled between 2011 and 2017–2018, outpacing the growth of total U.S. robot imports, consistent with the decline in Japanese robot export prices documented in Appendix Figure IA.2.1.

At the firm level, we ask whether adoption responded most where ex-ante exposure to the shock was highest. Appendix Table IA.3.2 interacts the yen with an indicator for a pre-period supplier relationship with Japan, defined as any 2007–2011 imports of *non-automation* goods from Japan so

[29] Since robots are themselves a component of CAPX, this faint positive drift in general investment is consistent with automation being the active margin rather than evidence of a broader confounding trend.

that the exposure has no mechanical overlap with the outcomes. Firms with a pre-existing Japan relationship increased automation adoption significantly more as the yen depreciated, and the effect survives restricting the sample to firms with some pre-period imports. [30]

Appendix Table IA.3.3 performs the analogous exercise with industry-level exposure to automation, using two ex-ante measures: the industry's pre-2011 automation import propensity, and the share of patents used by the industry that are automation patents, from Mann and Püttmann (2023), a measure of technological applicability external to the import data. The import-propensity interaction is positive and significant for all three outcomes, and the patent-based measure confirms the direction, significantly so for automation intensity. The response to the yen is therefore concentrated where geographic and technological exposure predict it, based on within-year cross-sectional variation rather than the aggregate episode alone.

Finally, the response is specific to automation equipment rather than to importing in general. Appendix Table IA.3.4 estimates the average effect of the yen on the probability that a firm imports in a given category: industrial robots (HS 847950), automation equipment broadly, and final consumer goods. The depreciation raises the probability of importing robots and of importing automation equipment, but leaves the probability of importing final goods unchanged. A general expansion of import demand, for example if firms were responding to broad US dollar strength, would move final goods as well. We thus conclude that the response to the yen appears in a way consistent with our proposed mechanism.

5.2.3 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

[30] We use this average differential rather than interacting the exposure with financial constraints: bilateral trade with Japan is a coarse proxy for exposure to Japanese technology, since automation systems bundle components from several countries. The model implies that the constrained-firm differential operates *through the balance sheet rather than the supplier network*, so there is no prediction that it should be stronger among Japan-connected firms. This is the larger reason that we don't use a more granular exposure measure: Firms in our model are differentially exposed based on their balance sheet and it wouldn't make sense to measure exposure through something like a link to Japan. Furthermore, since we are in large part capturing an *extensive margin* adoption effect, new adopters will not necessarily have any previous ties to Japan as a trade source.

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.[31] 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.[32]

Given the propensity scores, Panel A of Table 6 presents the results when re-estimating our baseline panel specification with inverse propensity weights (IPW). This method preserves all firm-year observations for which a propensity score can be estimated, but reweights them so that 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 6 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.[33]

Overall, the persistence of the constraint effects across both weighting and matching estimators

[31] 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.

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

[33] 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.

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 summarizes all observable factors that systematically distinguish adopters from non-adopters. We then implement two complementary tests reported in Appendix Table IA.3.5.

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 three columns, we adopt a complementary strategy that restricts comparisons to firms operating within the same industry–year cells by including industry $\times$ year fixed effects. This specification absorbs all common shocks at the industry–year level, so identification comes from comparisons among firms facing similar product-market conditions. The results remain 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 eventual adoption of firms that are inherently more likely to adopt, but by differential investment responses within the set of firms that are already viable adopters.

5.2.4 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.[34]

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.[35] 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

[34]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.

[35]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.

the rate elasticities: that more constrained firms will respond more aggressively to financing shocks along with price shocks.

We implement the empirical strategy described in equation 5.2, where the shock is the change in the one-year Treasury rate (lagged). We instrument this change using the orthogonalized monetary policy surprise series constructed by Bauer and Swanson (2023), which isolates unexpected policy-driven movements in short-term interest rates and purges information effects.

In the baseline specification in Table 7, Panel A we see that interest rate increases systematically reduce automation adoption. Higher Treasury rates significantly decrease automation across all measures. However, these aggregate effects mask substantial heterogeneity across firms as seen in Panel B of Table 7.[36]

TABLE 7 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 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

[36]The excluded instrument is highly relevant in the firm-level panel. The first-stage Kleibergen–Paap rk Wald statistic is well above the Stock–Yogo rule-of-thumb threshold of 10, reflecting the substantial precision gained from observing each annual monetary-policy surprise across thousands of firms. Such strong first-stage performance is typical in exposure-style designs, where identification arises from heterogeneous responses to a common shock rather than repeated time-series realizations of the shock. Our inferential concern is therefore valid inference with few aggregate shocks, in this case 13; accordingly, Table IA.3.7 checks robustness using procedures designed for small-T dependence.

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:

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

where the outcome variable y_{it} is a measure of the investment of firm i in year t in automation. $\hat{R}_{t-1}$ 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.

Figure 7 plots the impulse response functions implied by estimating equation 5.4, where the dependent variable is defined as the cumulative change in automation investment intensity between t-1 and $t+h$. Panel (a) presents the response for all manufacturing firms to a 25 basis point decrease in the 1-year Treasury rate, instrumented with the high-frequency with the orthogonalized monetary policy shock. The figure shows that automation investment rises on impact, with an increase of approximately 1.2 units per 1,000 assets in the year of the shock. The cumulative response gradually declines with the horizon, but remains positive through year 4, although statistical precision decreases at longer horizons.

Panels (b) and (c) decompose this aggregate response by firms' financial conditions. Panel (b) shows that firms with high cash constraints exhibit a substantially larger and more persistent cumulative response to monetary policy easing. For these firms, automation intensity increases by roughly 2 units on impact the year of the shock and remains significantly elevated the subsequent years, with the cumulative effect declining only gradually as the horizon increases. This pattern indicates that financing frictions amplify the transmission of monetary policy to automation investment.

In contrast, Panel (c) shows that firms with low cash constraints display a much smaller cumulative response. The estimated impulse responses are small and generally not statistically different from

zero, suggesting that monetary policy shocks have limited effects on automation investment among financially unconstrained firms.[37]

FIGURES 7 AND 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 8, 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) report the event-study responses of R&D spending for high- and low-cash-constraint firms, respectively. While R&D spending declines in the post-shock period for both high- and low-constraint firms, this pattern is best interpreted as reflecting re-timing or reallocation of innovative activity around a baseline level of product development, rather than a true, scalable interest-rate elasticity comparable to that observed for automation investment.

Alternative Measures of Shocks to Financing As an additional robustness check on our monetary policy results, Table IA.3.6 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

[37]The magnitudes in the IRF figures are not directly comparable to the regression coefficients in the tables, as they reflect cumulative changes in automation investment scaled to a 25 basis point monetary policy shock and estimated separately across samples.

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.

5.2.5 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 IA.4.1 and IA.4.2 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 8 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 5 and 7, with the relative price responsiveness now marginally larger and the relative interest rate responses slightly lower.

We can quantify the relative responses of high vs. low constrained firms, and furthermore compare across the price and interest rate shocks in a standardized manner. We interpret the two interaction coefficients of Column (1), Table 8, as the differential effects, while normalizing the average effect of each shock with the results in Panel A of Tables 5 and 7 respectively. In this way, the adoption probability response of constrained firms to the price shock is 2.5 times larger than that of unconstrained firms for the same shock. With respect to the monetary policy shock, the ratio is 1.6.[38] Across the 3 outcome measures this relative effect is quantitatively similar.[39] Section 6.3 shows that these two differentials, filtered through the model, recover the advance rate on robot hardware and its cross-firm gap.

5.2.6 Robustness

We subject the above regression results to several robustness tests. First, Table 9 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.

[38] See Footnote 27 for how these magnitudes are calculated. The average effects for the price shock and the monetary policy shock (MPS) are 0.225 and -0.09 respectively (Panel A of their respective Tables). The "differential" βs are 0.213 and -.045 (Table 8). Thus, the referenced calculation yields a differential effect of 2.5 for the price shock and 1.6 for the MPS, or a 56% larger relative effect following the price shock. The same procedure can be replicated for all outcomes.

[39] To reiterate, notice this is a statement about the cross-firm *differential* and not a claim that the price shock was the larger aggregate driver of adoption. Standardized by each shock's own identifying variation, the two differential responses are in fact of comparable magnitude (Section 6.3).

These should be similar to Linn and Weagley (2024), but are only available until 2015.[40] Using this text-based constraint measure, we find qualitatively similar results to those in Table 8. 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 9 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.3.8 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. 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,

[40] Because the Hoberg–Maksimovic (HM) financial-constraint measure is only available through 2015, we assign each firm its last observed HM status to all subsequent years through 2019. Firms that first enter Compustat after 2015 have no observed HM status and therefore cannot be classified using this measure.

potentially creating spurious correlations in our results.

To address this concern, Table 10 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.[41] 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 10 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.[42]

Third, a key question in the literature concerns which firms drive aggregate investment responses to macroeconomic shocks. Kroen et al. (2025) 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. 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 11.

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. (2025) and examine responses among the remaining smaller manufacturers. The results reveal that financially constrained laggards, rather than industry leaders,

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

[42] 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.

drive automation adoption in response to macroeconomic shocks.[43] In Panel B, we add several controls to the baseline specification. Namely, we include Tobin's Q (Peters and Taylor (2017))[44]; 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 11 ABOUT HERE

The internet appendix includes several other robustness tests that further reject alternative mechanisms. Table IA.3.9 excludes automotive sector firms (NAICS 336), a particularly robot-intensive industry, and confirms that our findings are not driven by this specific subsector. Table IA.3.10 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. We note that the rate channel carries half of the model's joint restriction, and it is not an economically marginal one: standardized by identifying variation, the constrained–unconstrained differential for the rate shock is comparable in magnitude to the price differential (Section 6.3), and the credit-spread and excess-bond-premium replications in Table IA.3.6 show the result is not specific to the Treasury-yield instrument.

[43] We don't view this as a conflict in mechanisms, but rather that Kroen et al. (2025)'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.

[44] 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.

6 Discussion and Model Calibration

6.1 A Counterfactual Exercise: Automation Adoption Without the Yen Depreciation

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 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.

To quantify the economic significance of the heterogeneous elasticities we estimate, we ask how many robots financially constrained firms would have adopted had the yen remained at its pre-Abenomics level. Specifically, we fix the exchange rate at its 2011 value for the 2012–2015 window, multiply the resulting change in the yen by the estimated coefficient on *High Cash Constraint* $\times$ *Yen* from Table 8, recompute IHS-transformed shipments for each constrained firm-year, and recover levels via the inverse hyperbolic sine. Because the simulation uses only the differential interaction coefficient, it isolates the additional adoption that the price decline induced among constrained firms relative to unconstrained firms, whose predicted imports are mechanically unchanged.

Under this counterfactual, financially constrained firms would have imported approximately 353 fewer robots over 2012–2015, a 17.3 percent shortfall relative to observed adoption. This aggregate effect masks substantial heterogeneity across industries, as shown in Figure 9. In several price-sensitive sectors the implied shortfall exceeds 20–30 percent (e.g. Chemical Manufacturing, NAICS 325; Other Manufacturing, NAICS 339). Part of this variation reflects differences in price 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.

FIGURE 9 ABOUT HERE

The 17.3 percent figure is also large relative to what our estimates imply for the financing channel when the two shocks are compared at realized magnitudes. An analogous counterfactual for a sizeable 100 basis point monetary policy shock implies a shortfall of only about 3-4 percent among constrained firms, and the actual high-frequency policy surprises in our sample are smaller than this benchmark. (A raw comparison of the two interaction coefficients in Table 8 is not informative, since one is per unit of log exchange rate and the other per percentage point of yield; the realized-shock comparison here and the standardized comparison in Section 6.3 put the two shocks on explicit footings.) 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.

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

The counterfactual exercises quantify how much the yen shock mattered; we now ask whether the *source* of the heterogeneous response is identifiable. A standard "user-cost plus wedge" specification—$k_i' = g(u_i(p_k, R), \omega_i)$ with a firm-specific friction ω_i rescaling user-cost responses—can always rationalize the fact that more constrained firms exhibit larger elasticities: for any joint pattern $\{\varepsilon_{k,p}, \varepsilon_{k,R}\}$, one can choose g and a mapping from observables to ω_i that fits it.

Our collateral model is more restrictive because it ties both elasticities to the same balance-sheet primitives, and single-primitive shifts produce distinct joint patterns. Proposition 3 provides the key diagnostic. A liquidity tightening ($c \downarrow$) uniquely generates (i) larger $|\varepsilon_{k,p}|$, (ii) larger $|\varepsilon_{k,R}|$, and (iii) a tighter financial marker (higher λ at the cap or higher θ^* in the interior). A reduction in non-pledgeable intangibles ($h \downarrow$) also amplifies both elasticities but *loosens* the financial margin, producing the opposite movement in the marker. No other single primitive produces all three movements simultaneously.

Two implications follow. First, the marginal adopters are low-c firms whose high effective cash shares

make them acutely sensitive to the sticker price net of collateralized borrowing. When robot prices fall, their per-unit cash outlay drops sharply, moving them across the adoption threshold—consistent with the 17 percent counterfactual shortfall documented above. Second, because the empirical pattern rules out h as the driver, the model's joint restriction on price and rate elasticities is not merely a convenient parameterization but a falsifiable prediction that the data confirm.

These findings speak to several strands of literature. They contribute to work 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. The results also connect to research on market structure and technological diffusion. If macro shocks disproportionately expand automation among constrained firms, they compress technological gaps; if they raise financing costs or input prices, they widen them. Macroeconomic policies and external shocks therefore shape not only the aggregate pace but also the distribution of technological upgrading.

Taken together, the evidence places our contribution at the intersection of financing constraints, collateral economics, and the diffusion of frontier production technologies. Exchange-rate–driven price shocks generate the largest constrained-firm adoption differentials—larger than those generated by interest-rate shocks—and the empirical patterns are consistent with a marginal adopter who is cash-poor rather than intangible-light. This 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.

6.3 From Elasticities to a Financing Primitive: Calibrating the Advance Rate

The diagnostic above uses the *pattern* of the two elasticities. We now push one step further and ask whether their *magnitudes*, filtered through the model, recover a financing primitive that can be checked against independent data. At the collateral cap, both shocks transmit through the invoice with a common factor: from $k' = c/M$, $\partial \ln k'/\partial \ln M = -1/(1+\kappa)$, while the shock-specific responses are $\partial \ln M/\partial \ln p_k = W_{p_k}$ and $\partial \ln M/\partial \ln R = -W_f\, \chi_R$ (Proposition 1). On the hardware component,

both magnitudes are governed by a single number: the *advance rate* of Corollary 1,

$$\varphi \equiv y\, m(\bar{\theta}), \tag{6.1}$$

the fraction of the hardware price a lender will finance, so that $W_{p_k} \propto (1-\varphi)$ (paid in cash) and $W_f \propto \varphi$ (financed). The following calibration recovers φ from two independent combinations of our estimates and checks both against firms' observed financing behavior.

The cross-firm gap in φ. The first combination uses only the amplification ratios that constitute the paper's headline heterogeneity: $E_C/E_U = 2.5$ for the price shock and 1.6 for the rate shock (see section 5.2.5 for the derivation using the average effects from Tables 5 and 7 and the differential effects from Table 8). To be clear, we use the notation "C" and "U" to differentiate firms, but in reality compare firms across *high* relative to *low* levels of constraints. In fact, the derivation below requires both sets of firms to be within the finance-binding region in the model, and we think this is the correct application given the empirical exercise simply differentiates firms by the magnitude of their constraints. In this environment, within a shock, the outcome scaling, the shock units, and χ_R are common to constrained and unconstrained firms and cancel, so these ratios and their ratios are unit-free:

$$\frac{(E_C/E_U)_{\text{price}}}{(E_C/E_U)_{\text{rate}}} = \frac{W^C_{p_k}/W^U_{p_k}}{W^C_f/W^U_f} = \frac{(1-\varphi_C)/\varphi_C}{(1-\varphi_U)/\varphi_U} = \frac{2.5}{1.6} = 1.56. \tag{6.2}$$

More constrained firms carry a roughly 56% higher cash-to-financed ratio on the pledgeable hardware. Anchoring the unconstrained advance rate at any plausible level, the implied cross-firm *gap* is remarkably stable: about 10 percentage points (Table 12, Panel A). Stated plainly, for the same piece of equipment a constrained firm can pledge less and must cover roughly 10 percentage points more of the hardware cost from internal funds—a larger down payment. This is the advance rate of the collateral literature: Rampini and Viswanathan (2010, 2013) make financing synonymous with collateral, so weaker firms pledge less; Rampini (2019) shows constrained agents rationally make larger down payments on durable assets; Kermani and Ma (2023) tie advance rates to assets' liquidation values; in macroeconomics it is the loan-to-value ratio of Kiyotaki and Moore (1997).

TABLE 12 ABOUT HERE

The level of φ_C. The second combination compares the two differentials *across* shocks, which requires standardizing the coefficients in Table 8. The model gives $|\varepsilon_{k',p_k}/\varepsilon_{k',R}| = (1-\varphi)/(\varphi\,|\chi_R|)$, so the standardized cross-shock ratio pins the level of φ given χ_R (Corollary 1). We standardize each interaction coefficient by the dispersion of its shock's identifying variation, set $\chi_R = -1$ as a benchmark (one-for-one pass-through of the policy rate into the lender's advance), and treat the estimated differentials as approximating constrained firms' own responses.[45] On a per-standard-deviation footing the two differentials are nearly equal across all three outcomes, and the implied level is $\varphi_C \approx 0.50$ (Table 12, Panel B): a constrained adopter finances about half of the hardware and pays the rest down in cash. The level is stable across outcomes and across standardizing by the interquartile range instead of the standard deviation (0.51 vs. 0.53), and moderately sensitive to the pass-through benchmark ($\varphi_C \in [0.36, 0.63]$ for $\chi_R \in [-1.5, -0.5]$).

Both outputs match non-targeted financing data. As a financing proxy, we use net new debt per dollar of capital expenditure, Δ(total debt)/capex (winsorized at the 5th/95th percentiles), among robot adopters split by our cash-constraint indicator. Unconstrained adopters finance 64 cents of each capex dollar with new debt; constrained adopters finance 52 cents (Table 12, Panel C). Two independent model outputs thus line up with two independent data moments, among the firms the model is actually about: the *level* ($\varphi_C \approx 0.50$ vs. measured 0.52) and the *gap* (≈ 10 pp vs. measured 12 pp)—neither of which was targeted in estimation.[46]

[45] The approximation is that "unconstrained" firms (in reality, less constrained) are essentially unresponsive, consistent with the decomposition in footnote 27, where E_U is small relative to E_C, and with the near-zero local-projection responses of unconstrained firms in Figure 7.

[46] As a caveat, Δdebt/capex is a net-flow measure rather than a contractual loan-to-value. Loan-level advance rates from secured credit facilities or equipment-finance loan-to-value schedules by borrower credit grade provide a more consistent check if the data is available. Furthermore, we concede there is no robot-specific advance-rate benchmark in existing data, so the level here is disciplined by the financing proxy found in the collateral literature (advance rates on new, redeployable machinery are high for strong credits, with weaker borrowers facing larger down payments).

7 Conclusion

This paper shows that the financing structure of a capital good shapes who adopts it. Contrary to the conventional view that financial constraints simply limit investment, we find that financially constrained firms are the *marginal* adopters of automation, responding more strongly than unconstrained firms to both lower equipment prices and lower financing costs. The difference is largest for capital-goods prices, indicating that price shocks have an especially important role in reallocating technology adoption across firms.

We rationalize these patterns with a collateral model in which the robot hardware is pledgeable but deploying it requires non-pledgeable, internally funded complementary investments. The friction binds not because the equipment cannot serve as collateral, but because putting it to work requires complementary spending that cannot. The model summarizes the financing friction in a single object—a per-unit cash invoice—and shows that firms facing a larger cash invoice become the marginal adopters when investment conditions improve. Although the model is estimated using only firms' adoption responses, it implies the advance rate lenders extend against robot hardware—near one-half and roughly ten percentage points lower for constrained firms—consistent with observed equipment-financing practices.

These findings have implications well beyond industrial robots. As production increasingly relies on software, customization, worker training, organizational change, and other complementary investments bundled with physical equipment, the binding constraint on technology adoption lies in financing the parts a lender cannot repossess. Whether a firm can finance the complementary investments that make a technology productive—not simply borrow against its tangible components—is what increasingly determines who adopts and who is left behind.

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 u.s. 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*.

Barbiero, O., Sheremirov, V., Silva, A., and Stein, H. (2026). Trade wars on the income statement: The effects of tariffs on firms. Available at SSRN 6844063.

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 firm investment. *Journal of Financial Economics*,

174:104183.

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 u.s. manufacturing. *AEA Papers and Proceedings*, 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.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annu. Rev. Econ.*, 1(1):451–488.

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.

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.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of economics and statistics*, 80(4):549–560.

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 July 2026.

Falato, A., Kadyrzhanova, D., Sim, J., and Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the u.s. 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. (2026). Firm balance sheet liquidity, monetary policy shocks, and investment dynamics. *Journal of Political Economy*. forthcoming.

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

Kermani, A. and Ma, Y. (2023). Asset specificity of nonfinancial firms. *The Quarterly Journal of Economics*, 138(1):205–264.

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. (2025). Falling rates and rising superstars. NBER Working Paper.

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. (2022). Low interest rates, market power, and productivity growth. *Econometrica*, 91(3):1015–1053.

Mann, K. and Püttmann, L. (2023). Benign effects of automation: New evidence from patent texts. *Review of Economics and Statistics*, 105(3):562–579.

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.

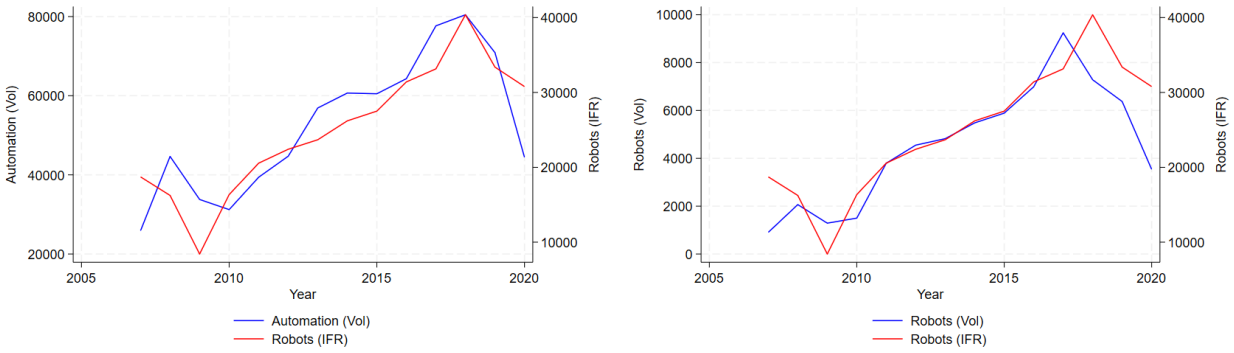
Rampini, A. A. (2019). Financing durable assets. *American Economic Review*, 109(2):664–701.

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 2: 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.

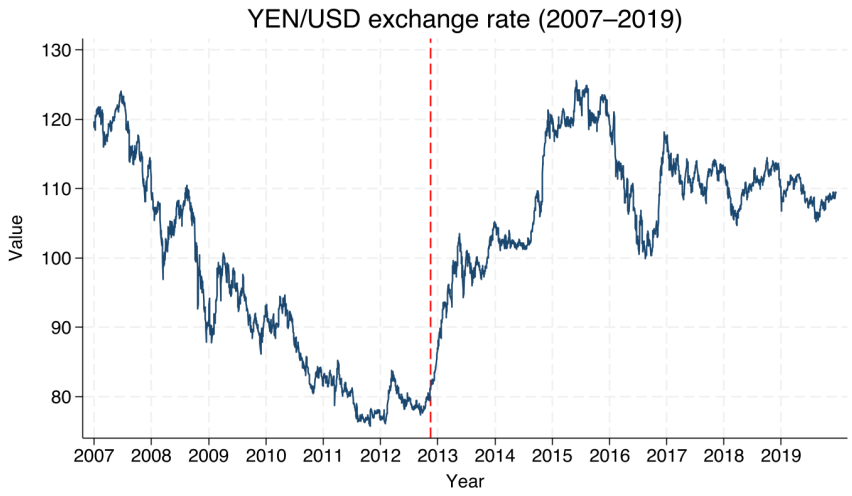


Figure 3: Japanese Market Share and the Yen-USD Exchange Rate

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.

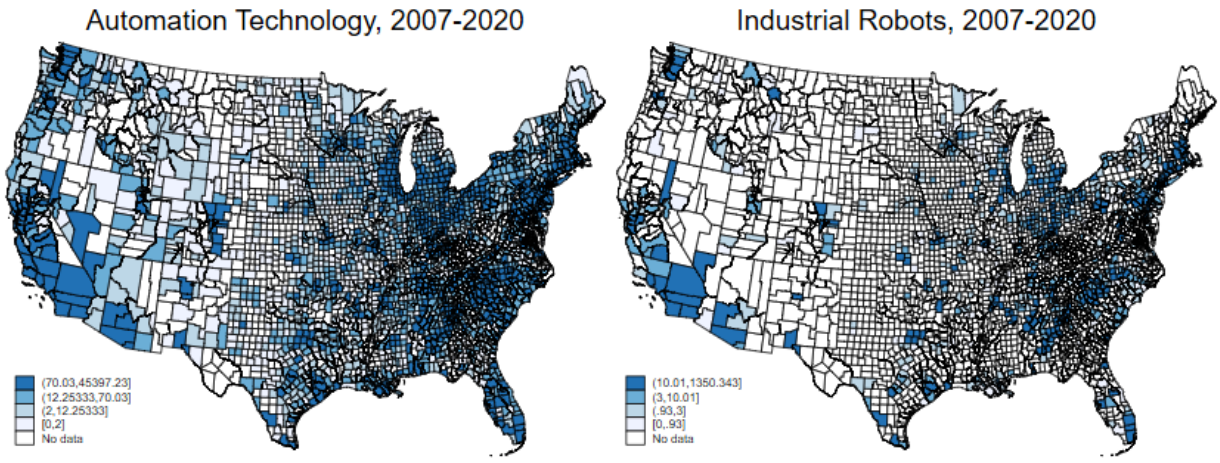
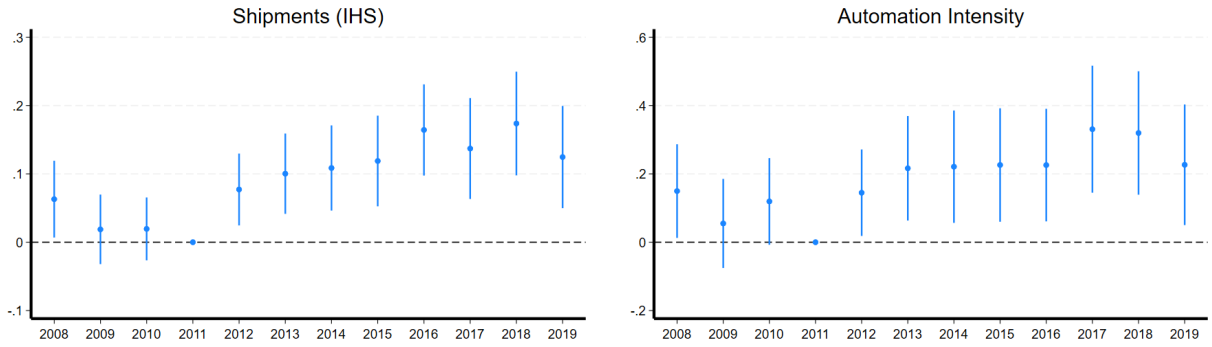


Figure 4: Geographic Distribution of Automation and Robots, 2007 - 2020
These figures shows the maps with all transactions of automation technologies *(Panel A)* and industrial robots imported *(Panel B)* in volume in 2007-2019 by county, excluding imports from logistic firms.

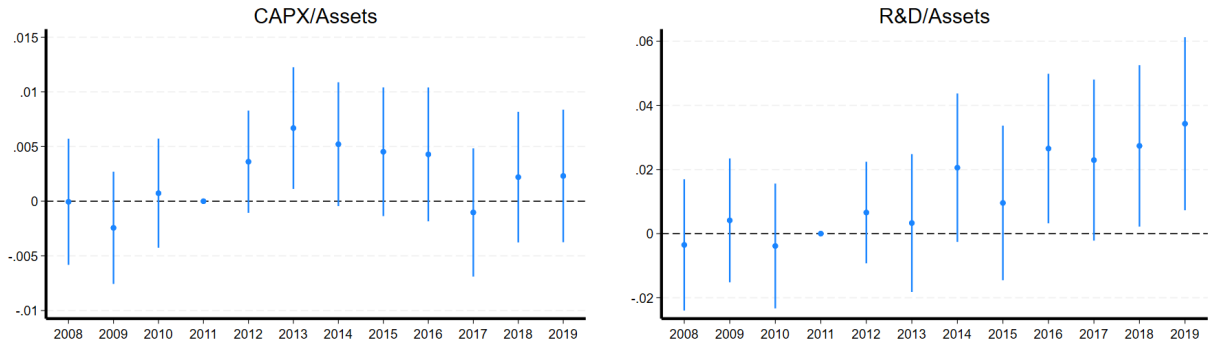


(a) Number of shipments

(b) Automation intensity

Figure 5: 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.

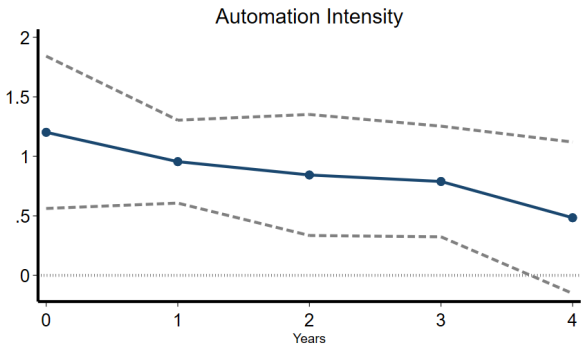


(a) Capital expenditure by assets

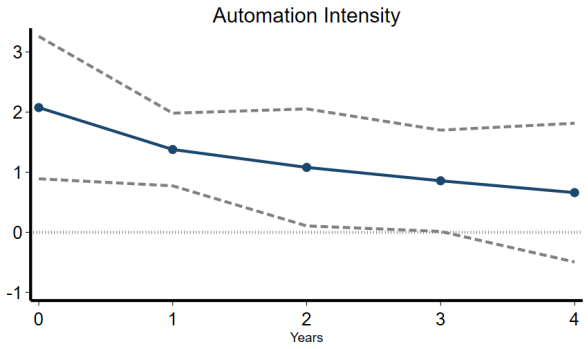
(b) R&D by assets

Figure 6: Differential Investment Response to YEN-USD Depreciation: CAPX and R&D by Cash Constraints

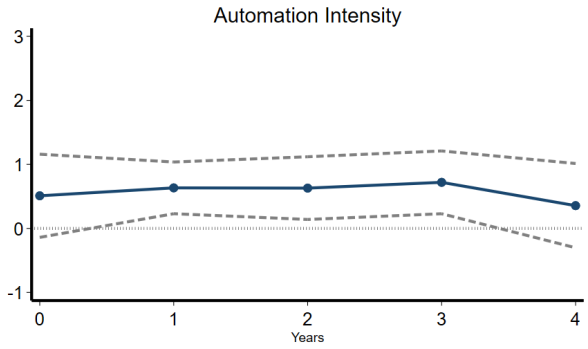
This figure plots the dynamic response of capital expenditure (CAPX) scaled by assets in Panel (a), and of R&D spending scaled by assets in Panel (b), to the Japanese Yen depreciation. The sample includes all manufacturing firms from 2007 to 2019. We measure firm exposure to the exchange rate shock as 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



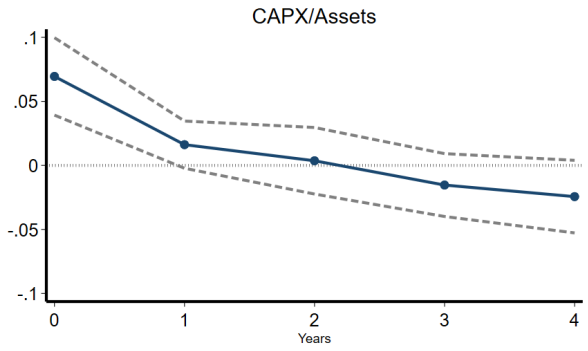
(b) Firms with high cash constraints



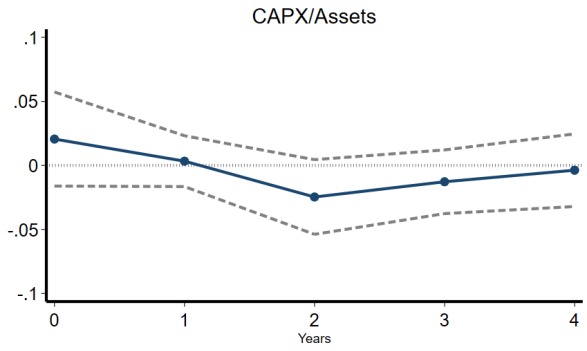
(c) Firms with low cash constraints

Figure 7: Differential Automation Response to Monetary Policy Shocks: By Cash Constraints

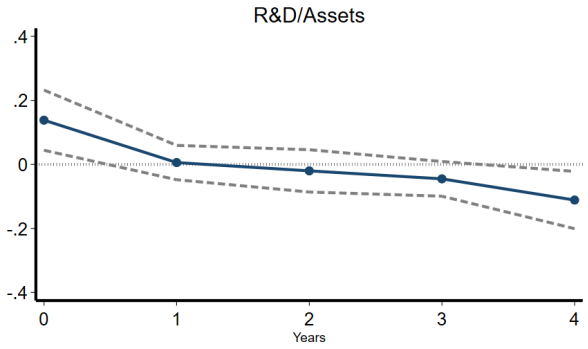
This figure plots the dynamic response of investment in automation technology (scaled by assets) to a 25bp decrease in the 1-year Treasury rate estimated using equation 5.4, instrumented by cumulative high-frequency monetary policy shocks. The sample includes all manufacturing firms from 2007 to 2019. 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.



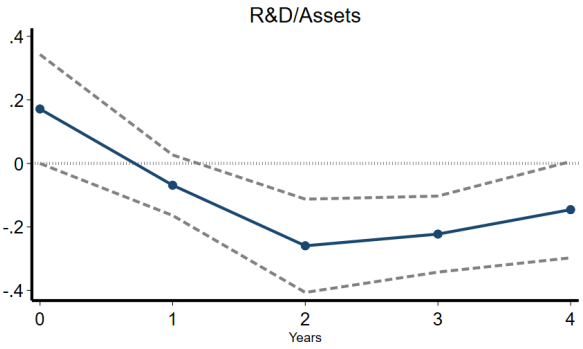
(a) Firms with high cash constraints



(b) Firms with low cash constraints



(c) Firms with high cash constraints



(d) Firms with low cash constraints

Figure 8: 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 25bp 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.

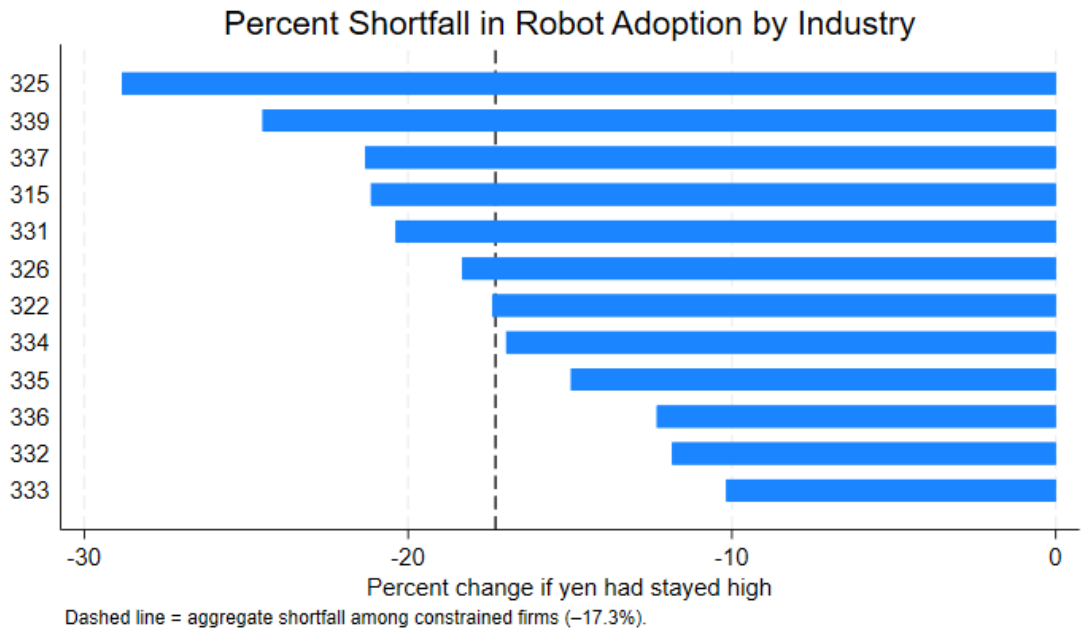


Figure 9: Counterfactual: Automation Without the YEN Depreciation, by 3-digit Industry

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 2: 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 shows the top six countries supplying industrial robots; both panels cover the 2007–2012 period and report each country's share of total import volume and total shipments, while Panel B presents the top five countries by share of total automation technology exports to the U.S. Panel C identifies the top ten individual companies supplying automation technologies and industrial robots to U.S. firms over the full 2007–2019 sample, 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. Country of origin is assigned from the shipment's place of receipt, mapped to a country via UN/LOCODE and defaulting to the shipment-origin country when unresolved (following Barbiero et al. (2026)).

Panel A: Robots by Country of Origin

Country	Volume (%)	Shipments (%)
Japan	53.13	39.38
Germany	13.62	12.91
China	6.47	12.70
South Korea	4.93	4.37
Italy	3.16	3.10
Belgium	3.08	4.90
Others	15.62	22.64
Total	*100*	*100*

Panel B: Automation by Country of Origin

Country	Volume (%)	Shipments (%)
Japan	20.77	22.22
China	20.21	23.00
Germany	14.41	10.54
South Korea	10.53	7.06
Taiwan	8.90	10.03
Others	25.18	27.15
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 3: 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 Industrial Robots

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

Panel C: Top Importers of Automation

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

Table 4: 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 5: 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 6: Automation and the Exchange Rate Shock: Propensity Score Matching
This table examines specifications with propensity scores that 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 7: 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

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

Table 8: 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

Table 9: 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

Table 10: 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)
$High\ CashCon_i \times Tyield_{t-1}$	-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 11: 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 Kroen et al. (2025). In Panel B, we repeat the specification in Table 8 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

Table 12: Calibrating the Advance Rate φ from the Estimated Elasticities

Panel A applies Equation (6.2) with $(E_C/E_U)_{\text{price}} = 2.5$ and $(E_C/E_U)_{\text{rate}} = 1.6$ (adoption outcome; footnote 27). Panel B standardizes each Table 8 interaction coefficient by the dispersion of its shock's identifying variation (for the rate, the component of the one-year yield explained by cumulated surprises, s.d. ≈ 0.66) and inverts $|\varepsilon_{p_k}/\varepsilon_R| = (1-\varphi)/(\varphi|\chi_R|)$ at $\chi_R = -1$; using the interquartile range instead of the standard deviation moves φ_C from 0.51 to 0.53, and $\varphi_C \in [0.36, 0.63]$ for $\chi_R \in [-1.5, -0.5]$. Panel C reports net new debt per dollar of capital expenditure, winsorized at the 5th/95th percentiles; the adopter row is the model-relevant comparison.

Panel A: The cross-firm gap from the amplification ratios (Eq. 6.2)

φ_U (anchor)	φ_C (implied)	Gap $\varphi_U - \varphi_C$
0.40	0.30	10 pp
0.50	0.39	11 pp
0.60	0.49	11 pp
0.70	0.60	10 pp

Panel B: The level from the standardized cross-shock ratio ($\chi_R = -1$)

Outcome	Per-σ yen differential	Per-σ rate differential	Implied φ_C
Automate dummy	0.029	0.030	0.51
Shipments (IHS)	0.050	0.054	0.52
Automation intensity	0.092	0.093	0.50

Panel C: Non-targeted financing moments, Δ(total debt)/capex

Sample / constraint split	Unconstrained	Constrained	Gap
Robot adopters, cash constraint	0.64	0.52	12 pp
All firms, cash constraint	0.93	0.59	34 pp
All firms, Linn–Weagley debt constraint	0.67	0.39	28 pp

Table A1: 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.2.1) 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/USD 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(CNY/USD 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(Nominal Emerging Market Economies U.S. Dollar Index divided by 100). 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
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

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 provides an overview of related work. Section IA.2 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.3 reports additional robustness tests and alternative specifications discussed in Section 5 of the paper. Section IA.4 presents proofs to the main propositions as well as theoretical extensions.

IA.1 Relation to Closely Related Work

This appendix situates our contribution relative to closely related work along three dimensions central to identifying mechanisms in investment responses: the type of shock studied, the measurement of investment activity, and the source of heterogeneity across firms. As Table IA.1.1 summarizes, most existing work varies one dimension at a time—studying financing shocks using broad investment measures, or documenting heterogeneous responses without capital-good-specific price variation.

Our contribution is to vary these dimensions jointly. We are, to our knowledge, the first to observe both capital-input price and interest-rate shocks on the same adoption margin. This joint observation is essential: absent transaction-level prices, one cannot separate capital-input price effects from financing-cost effects within the same investment category; absent variation in both shocks, different financing mechanisms – collateral constraints, liquidity constraints, complementary intangible investment – generate observationally equivalent responses. Our setting allows us to decompose user cost, show that the same firms are marginal for both shocks, and distinguish these mechanisms through a testable cross-equation restriction linking price and rate elasticities.

Note: The table is selective and intended to clarify identifying variation and measurement, not to rank contributions.

Table IA.1.1: Positioning relative to closely related work.

Paper	**Shock**	**Outcome**	**Key heterogeneity**	**Mechanism / testable restriction**	**Contribution relative to this paper**
This paper	Capital-input prices *and* interest rates	Transaction-level automation adoption (BoL imports linked to Compustat)	Collateral & liquidity tightness	Joint price–rate restriction tied to a single cash-invoice margin	Introduces transaction-level adoption elasticities to capital-input prices and borrowing costs measured on the same margin, and decomposes the relative importance of price versus rate channels..
Almeida and Campello (2007)	Credit supply	CAPX and broad investment measures	Financial constraints	Collateral and liquidity channels; single-shock comparative statics	*Their focus:* Foundational constraints–investment logic. *What we add:* Transaction-level adoption, capital-input price variation, and a joint price–rate restriction on the same margin.

Continued on next page.

Table IA.1 continued.

Paper	Shock	Outcome	Key heterogeneity	Mechanism / testable restriction	Contribution relative to this paper
Rampini and Viswanathan (2010)	Financing wedge	CAPX with endogenous collateral	Collateral scarcity	Collateral constraint with endogenous encumbrance	*Their focus:* Theory of collateral scarcity. *What we add:* Invoice-based adoption representation yielding closed-form price and rate elasticities and a directly testable cross-equation restriction.
Ottonello and Winberry (2020)	Monetary policy (rates)	Aggregate and firm investment	Financial conditions	Amplification of rate shocks; single-shock design	*Their focus:* Heterogeneous responses to rate shocks. *What we add:* Capital-input price variation and a joint price–rate restriction distinguishing mechanisms via cross-shock comovement.
Catherine et al. (2022)	Collateral-value shocks	Firm investment	Collateral values	Collateral revaluation channel	*Their focus:* Investment responses to collateral revaluations. *What we add:* Capital-input price shocks and a test of joint price–rate responses on a single adoption margin.
Döttling and Ratnovski (2023)	Monetary policy (rates)	Tangible vs. intangible investment	Asset intangibility	Single-shock asset-type comparison	*Their focus:* Heterogeneous monetary transmission across asset classes.[47] *What we add:* Joint price and rate shocks on one adoption margin and a cross-equation restriction unavailable in single-shock designs.

[47] Cloyne et al. (2023), Jeenas (2026), and Drechsel (2023) provide complementary evidence on heterogeneous monetary transmission; these designs do not combine capital-input price variation with a joint price–rate restriction on a single adoption margin.

IA.2 Sample Construction and Summary Statistics

Table IA.2.1: 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.2.2: Summary Statistics of Key Variables

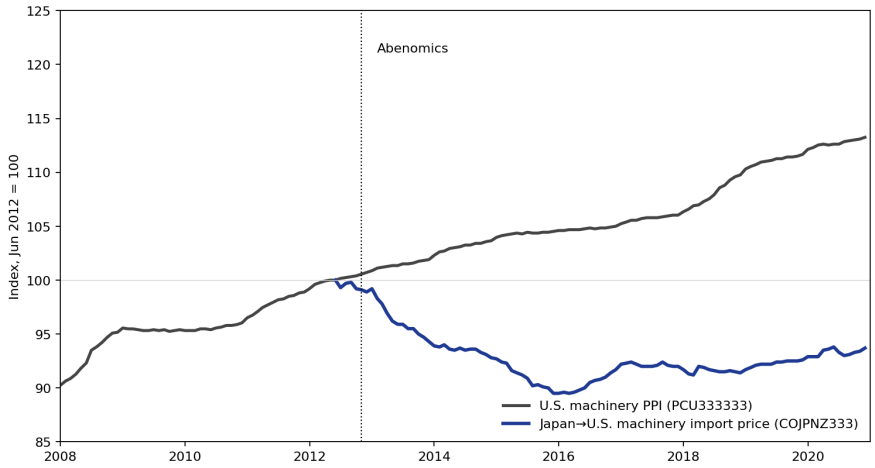
Variable	N	sd	mean	p25	p50	p75
Automate	21,974	0.359	0.152	0	0	0
Shipments (IHS)	21,568	0.587	0.208	0	0	0
Automation Intensity	21,635	1.299	0.300	0	0	0
Shipments Stock (IHS)	21,974	1.349	0.694	0	0	0.882
Volume (IHS)	21,586	0.686	0.214	0	0	0
YEN-USD	28,533	0.133	0.010	-0.065	0.056	0.099
Tyield	28,533	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.2.3: 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).

Figure IA.2.1: U.S. Machinery Producer Prices vs. Japan-Imported Machinery Prices



The figure plots two monthly price indices, both normalized to June 2012 = 100 (the natural base of the import-price series and the month immediately preceding the onset of Abenomics). The black line is the BLS Producer Price Index by Industry for Machinery Manufacturing, NAICS 333 (FRED series `PCU333333`), capturing the price of machinery produced and sold domestically in the United States. The blue line is the BLS Import Price Index for U.S. imports of NAICS 333 machinery originating in Japan (FRED series `COJPNZ333`), the most disaggregated country-by-product cell BLS publishes. The Japanese-specific series is only available starting in 2012. The vertical dotted line marks November 2012, the start of Japan's Abenomics monetary easing.

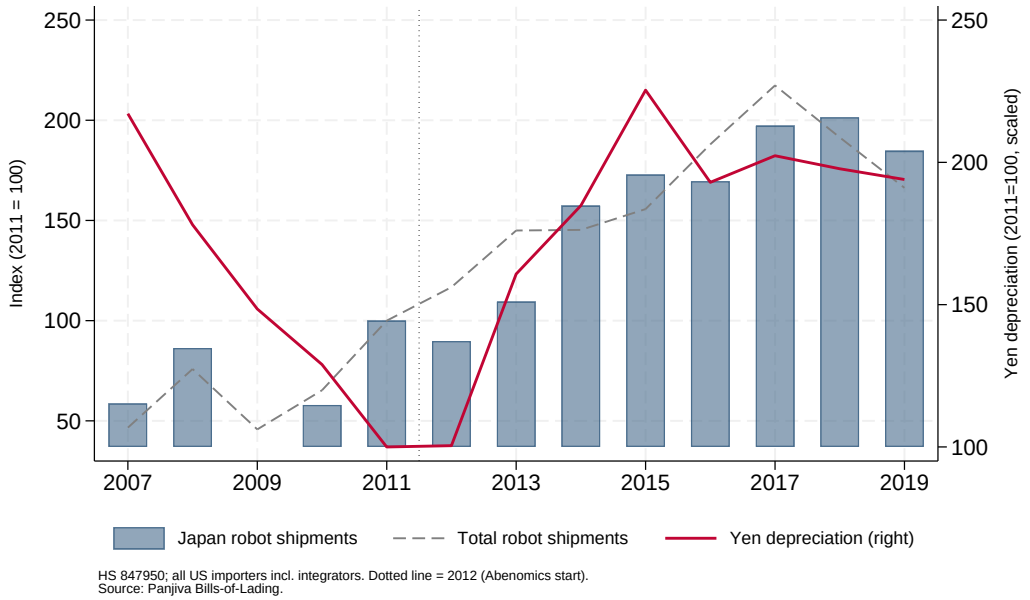
IA.3 Robustness Tests

Table IA.3.1: 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

Figure IA.3.1: Japan-origin robot imports and the yen



Notes: Industrial robot (HS 847950) import shipments into the U.S., indexed to 2011 = 100, for Japan-origin shipments and for all origins, plotted against the yen/dollar exchange rate. Origin assigned by the place-of-receipt method (Data Appendix). Full Panjiva import universe, including integrators and distributors. The dotted vertical line marks 2012, the start of the yen depreciation.

Table IA.3.2: Automation and the Exchange Rate Shock Robustness: Pre-Period Japan Supplier Relationships

This table examines whether the response of automation adoption to Japanese Yen depreciation is concentrated among firms with pre-existing supplier relationships with Japan. *Japan Rel* is an indicator variable equal to one if the firm imported any *non-automation* goods from Japan during 2007-2011; non-automation imports are used so that the exposure has no mechanical overlap with the outcome variables. The dependent variables and *Yen-USD* are identical to the baseline tests. Columns 1-3 use all manufacturing firms; Columns 4-6 restrict the sample to firms with some non-automation imports during 2007-2011, so that the coefficient does not reflect the contrast between importers and non-importers. All regressions include year and firm fixed effects, and standard errors (in parenthesis) are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
	All Manufacturing			Pre-Period Importers		
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$Japan\ Rel_i \times Yen - USD_t$	0.257***	0.635***	0.347*	0.189***	0.533***	0.136
	(0.051)	(0.094)	(0.208)	(0.056)	(0.101)	(0.231)
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	12,854	12,457	12,530
adj. R-sq	0.492	0.513	0.354	0.452	0.490	0.356

Table IA.3.3: Automation and the Exchange Rate Shock Robustness: Industry Automation Exposure

This table examines whether automation adoption responds most strongly to Japanese Yen depreciation in the industries most exposed to automation. *Ind Exposure* is one of two ex-ante, time-invariant industry exposure measures, standardized within manufacturing. In Columns 1-3 it is the industry (NAICS-3) share of firm-years importing any automation equipment during 2007-2011 (import propensity); because this measure is built from the same import data as the outcomes, it should be read as an industry-level proxy rather than an exogenous shifter. In Columns 4-6 it is the share of patents in use by the industry (SIC-2) that are automation patents over 2000-2011, from Mann and Püttmann (2023), a measure of technological applicability that is external to the import data. The dependent variable and *Yen-USD* measures are the same as the baseline analysis. All regressions include year and firm fixed effects, and standard errors (in parenthesis) are clustered at the firm level. The sample is restricted to Manufacturing firms. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
	Import-Propensity Exposure			Patent-Based Exposure		
	Automate Dummy	Shipments (IHS)	Automation Intensity	Automate Dummy	Shipments (IHS)	Automation Intensity
$Ind\ Exposure_j \times Yen - USD_t$	0.061***	0.154***	0.304***	0.018	0.053	0.136*
	(0.020)	(0.040)	(0.105)	(0.018)	(0.034)	(0.083)
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	21,637	21,230	21,295
adj. R-sq	0.491	0.511	0.355	0.490	0.510	0.354

Table IA.3.4: Automation and the Exchange Rate Shock Robustness: Yen Effect Across Import Categories

The dependent variable is an indicator equal to one if the firm imports at least once in the category that year: industrial robots (HS 847950); automation equipment (Acemoglu–Restrepo classification — our baseline Automate Dummy specification); final consumer goods (BEC consumption-goods classification). All specifications include firm fixed effects, firm controls (lagged log assets, log age, leverage, ROA), and macro controls (lagged real GDP growth, CPI inflation, industrial production). Standard errors (in parentheses) are clustered at the firm level. The sample is restricted to Manufacturing firms. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)	(3)
	Industrial Robots	Automation	Final Goods
$Yen\text{-}USD_t$	0.041**	0.225***	0.023
	(0.017)	(0.044)	(0.040)
FE	Firm	Firm	Firm
Firm Controls	Yes	Yes	Yes
N	21,637	21,637	21,637
adj. R-sq	0.150	0.489	0.724

Table IA.3.5: 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.3.6: 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

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

Table IA.3.7: Automation Response to Combined Exchange Rate and Monetary Policy Shocks: Robustness with Alternative Clustering

This table examines the joint effects of exchange rate and monetary policy shocks on firm-level automation investment, exploring alternative inference methods relative to Table 8. Each column follows exactly the same specification as that table, including firm and year FEs. The standard errors reported in parentheses are Driscoll and Kraay (1998) standard errors robust to general forms of cross-sectional and temporal dependence (up to 3 lags). This approach accounts for the potential cross-sectional correlation in error terms induced by applying aggregate time-series shocks to a panel of firms. Furthermore, the bottom rows report p-values from alternative clustering. We apply the Wild Cluster Bootstrap (WCB) clustered at the year level to account for the small number of aggregate shock periods. For the exogenous exchange rate shock ($YEN - USD$), WCB p-values are based on the Wald test statistic using Webb weights. For the endogenous interest rate shock ($Tyield$), we report p-values from the Anderson-Rubin (AR) test. Columns 1-3 define constraints using the High Cash Constraint index, while Columns 4-6 use the High Financial Constraint (Debt) index. Significance levels for the Driscoll-Kraay standard errors are indicated by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The sample is restricted to manufacturing firms from 2007-2019.

	(1)	(2)	(3)	(4)	(5)	(6)
	Automate	Shipments (IHS)	Automation Intensity	Automate	Shipments (IHS)	Automation Intensity
$High\ CashCon_i \times YEN - USD_t$	0.213**	0.369**	0.680**			
	(0.083)	(0.157)	(0.298)			
$High\ CashCon_i \times Tyield_{t-1}$	-0.045**	-0.082**	-0.141**			
	(0.018)	(0.031)	(0.054)			
$High\ FC_i \times YEN - USD_t$				0.178***	0.251**	0.428***
				(0.048)	(0.097)	(0.115)
$High\ FC_i \times Tyield_{t-1}$				-0.042***	-0.074**	-0.044
				(0.012)	(0.030)	(0.044)
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
Time Cluster Bootstrap pvalue (Yen)	0.01	0.02	0.02	0.00	0.04	0.08
Time Cluster Bootstrap pvalue (Rate)	0.04	0.00	0.01	0.00	0.06	0.69

Table IA.3.8: 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

Table IA.3.9: 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

Table IA.3.10: 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

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.4 Theoretical Appendix

This Appendix is organized as follows. Section IA.4.1 derives the elasticities for cash-rich firms from optimality conditions. To derive the elasticities for cash-poor firms, we first derive their cash identity in Section IA.4.2. For these firms, the level of investment is determined by borrowing the maximum allowed and investing in automation k' as much as allowed by their budget constraints. Section IA.4.3 derives the cap–regime elasticities ($\theta = \bar{\theta}$); Sections IA.4.4 derive and sign the interior–regime elasticities ($0 < \theta < \bar{\theta}$). Section IA.4.5 derives the comparative statics. Sections IA.4.6 provide results on the shadow value λ associated with the finance-binding constraints.

IA.4.1 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.4.2 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:

$$c = \big(p_k + h - y\,m(\theta)\,p_k\big)k' + \tfrac{\tau}{2}(k')^2 \tag{IA.4.1}$$

$$= k' \underbrace{\Big[p_k\big(1 - y\,m(\theta)\big) + h + \tfrac{\tau}{2}k'\Big]}_{M}, \tag{IA.4.2}$$

Define the *internal invoice per unit* M as the own-cash requirement accompanying each robot purchase:

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

The fundamental cash identity equates total required internal funds to available cash, $k'M = c$.[48] 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.

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.

There are two subcases of cash-poor firms that are similar but give rise to slightly different algebraic analysis below.

[48] For notional simplicity we will often write M and y without their arguments, keeping in mind that $M = M(k', \theta; R, \pi, \rho)$ and $y = y(R, \pi, \rho; r_c)$.

Cash-poor firms that are "at the cap" are those whose encumbrance 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 IA.4.1 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 $\tau > 0$.

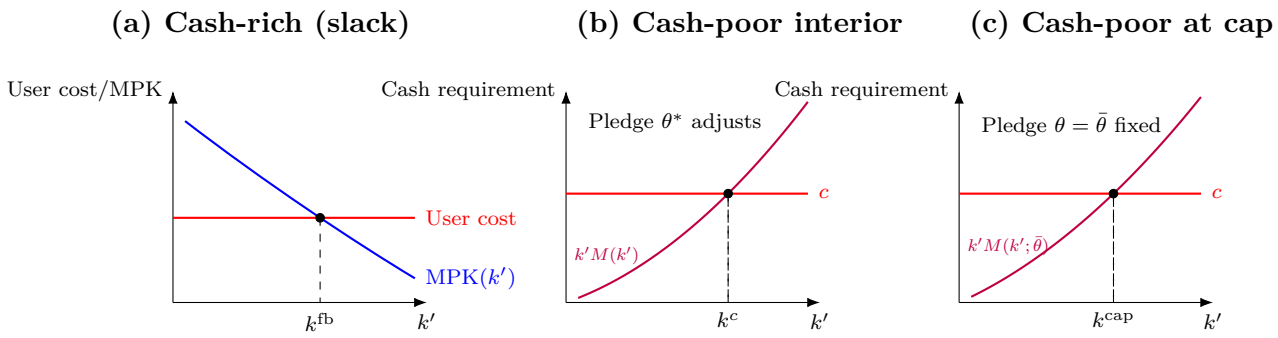


Figure IA.4.1: Three regimes. **(a)** Cash-rich (financing slack): choose k' where MPK(k') equals user cost. **(b)** Cash-poor, interior pledge: financing binds; k' solves $k'\,M(k';\theta^*) = c$. **(c)** Cash-poor at the cap: financing binds and $\theta = \bar{\theta}$, so k' solves $k'\,M(k';\bar{\theta}) = c$.

IA.4.3 Investment Elasticities for Cash-poor Firms in the Cap: Proposition 1

Proof. Consider the "Cap Regime", where the firm's borrowing capacity is exhausted in equilibrium ($\theta^* = \bar{\theta}$). From $k'M = c$ 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 $M = p_k\big(1 - y\,m(\bar{\theta})\big) + h + \frac{\tau}{2}k'$:

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

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

Third, substitute $d\ln M$ from the second step into the total derivative from the first step:

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

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).[49] By the same steps (using $\partial\ln M/\partial\ln R = -W_f\,\chi_R$), we obtain

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

[49] Notice that the sign follows from the following over-advance vs. hardware assumption: $y\,m(\theta) \le 1$ for all $\theta \in [0,\bar{\theta}]$. Lenders never advance more than the robot sticker cost per unit of effective collateral. With imperfect recovery ($\rho < 1$), enforcement frictions ($m(\bar{\theta}) < 1$), and material default risk ($\pi < 1$), the advance rate y is bounded well below $1/m(\bar{\theta})$. Intuitively, lenders facing default and recovery risk will not advance more than the purchase price of the collateral.

□

IA.4.4 Investment Elasticities for Cash-poor Firms with Interior Pledge: Proposition 2

Proof. We develop the proof of Proposition 2 in three subsections for clarity. Subsection IA.4.4.1 derive the pledge elasticities, used in the derivation of the price elasticity in subsection IA.4.4.2 and rate elasticity in subsection IA.4.4.3.

IA.4.4.1 Pledge elasticities: definition and origin

In the interior regime ($0 < \theta^* < \bar{\theta}$) with binding finance, the pledge choice trades off (i) *relaxing* borrowing capacity through higher effective pledgeability $m(\theta)$ and (ii) *reducing* equity's payoff in default because equity forfeits the legal encumbrance share θ while creditors recover only $m(\theta) \leq \theta$ (the residual is deadweight loss).

Working in the finance–binding region, we can write the reduced Lagrangian as

$$\mathcal{L}(k', \theta, \lambda) = U(k', \theta) + \lambda\big[c - k' M(k', \theta)\big],$$

where the invoice is $M(k', \theta) = p_k\big(1 + \frac{h}{p_k} - y\, m(\theta)\big) + \frac{\tau}{2}k'$ and the reduced-form objective (after imposing competitive pricing so b is neutral in PV) is

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

Hence

$$U_\theta(k', \theta) = -\frac{1-\pi}{R}\rho\, p_k k'\,\big(1 - m'(\theta)\big), \qquad M_\theta(k', \theta) = -p_k\, y\, m'(\theta).$$

The interior first–order condition for pledge is therefore

$$U_\theta(k', \theta^*) - \lambda\, k'\, M_\theta(k', \theta^*) = 0 \quad \Longleftrightarrow \quad \boxed{\lambda\, y\, m'(\theta^*) = \frac{1-\pi}{R}\rho\,\big(1 - m'(\theta^*)\big)}. \tag{IA.4.3}$$

Pledge elasticities. 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 condition (IA.4.3).

Price shock. Holding (R, π, ρ) fixed (so y is fixed), log–differentiate (IA.4.3):

$$\ln\lambda + \ln y + \ln m'(\theta^*) = \text{const.} + \ln\big(1 - m'(\theta^*)\big).$$

Differentiating with respect to $\ln p_k$:

$$d\ln\lambda + d\ln m'(\theta^*) = d\ln\big(1 - m'(\theta^*)\big),$$

where

$$d\ln m'(\theta^*) = \frac{\theta^* m''(\theta^*)}{m'(\theta^*)}\, d\ln\theta^*, \qquad d\ln\big(1 - m'(\theta^*)\big) = -\frac{\theta^* m''(\theta^*)}{1 - m'(\theta^*)}\, d\ln\theta^*.$$

Substituting and collecting terms:

$$d\ln\lambda = -\frac{\theta^* m''(\theta^*)}{m'(\theta^*)(1 - m'(\theta^*))}\, d\ln\theta^*.$$

Hence

$$\eta_{\theta,p} = \frac{\partial \ln \theta^*}{\partial \ln p_k} = -\frac{m'(\theta^*)\big(1-m'(\theta^*)\big)}{\theta^* m''(\theta^*)} \frac{\partial \ln \lambda}{\partial \ln p_k}. \tag{IA.4.4}$$

Since $m''(\cdot) < 0$, any monotonicity condition implying $\partial \ln \lambda / \partial \ln p_k > 0$ yields $\eta_{\theta,p} > 0$ whenever $0 < m'(\theta^*) < 1$ (which holds if $m(\theta) \le \theta$ and m is concave).

Rate shock. Log–differentiating (IA.4.3) with respect to R, and using $\chi_R \equiv \partial \ln y / \partial \ln R$:

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

Collecting the $\eta_{\theta,R}$ terms:

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

so

$$\eta_{\theta,R} = -\Big(1 + \frac{\partial \ln \lambda}{\partial \ln R} + \chi_R\Big) \frac{m'(\theta^*)\big(1-m'(\theta^*)\big)}{\theta^* m''(\theta^*)}. \tag{IA.4.5}$$

Without further structure, $\eta_{\theta,R}$ is not signed from primitives alone; below we rely on Assumption (A8), which directly restricts the combination that enters the rate elasticity.

IA.4.4.2 Price elasticity with interior pledge

In the interior regime ($0 < \theta^* < \bar{\theta}$) with binding finance, the cash identity $k'M = c$ implies $d\ln k' + d\ln M = 0$. Allowing θ to adjust endogenously,

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

Using $d\ln\theta = \eta_{\theta,p}\, d\ln p_k$ (with $\eta_{\theta,p}$ defined in §IA.4.4.1), we obtain

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

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

$$(1+\kappa)\, d\ln k' = -\big(W_{p_k} - W_\theta\, \eta_{\theta,p}\big)\, d\ln p_k,$$

so the interior-regime price elasticity is

$$\boxed{\varepsilon^{\text{int}}_{k',p_k} \equiv \frac{\partial \ln k'}{\partial \ln p_k} = -\frac{W_{p_k} - W_\theta\, \eta_{\theta,p}}{1+\kappa}.}$$

We assume incomplete pledge adjustment so that 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: $0 \le W_\theta \eta_{\theta,p} < W_{p_k}, \qquad W_\theta \eta_{\theta,R} + W_f \chi_R < 0$. This ensures that the numerator is positive and $\varepsilon^{\text{int}}_{k',p_k} < 0$.

IA.4.4.3 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

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

Table IA.4.1: 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 ($\pi \uparrow$)	$\downarrow$	$\downarrow$	$\uparrow$	fixed at $\bar{\theta}$
Higher Recovery ($\rho \uparrow$)	$\downarrow$	$\downarrow$	$\uparrow$	fixed at $\bar{\theta}$

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

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

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

$$\boxed{\varepsilon^{\text{int}}_{k',R} \equiv \frac{\partial \ln k'}{\partial \ln R} = \frac{W_f\,\chi_R + W_\theta\,\eta_{\theta,R}}{1+\kappa}.} \tag{IA.4.6}$$

Under Assumption (A7), $W_\theta \eta_{\theta,R} + W_f \chi_R < 0$, so the numerator in (IA.4.6) is negative and hence $\varepsilon^{\text{int}}_{k',R} < 0$.

□

IA.4.5 Comparative Statics

Tables IA.4.1 and IA.4.2 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.[50]

In the cap–regime changes to the non-pledgeable setup costs h and installation frictions τ requires that $\kappa = \frac{\tau c}{2M^2} < 1$, which ensures that the exposure effect dominates the installation-cost brake and therefore that the arrows in Table IA.4.1 are locally signed.

IA.4.6 Monotone Tightness and the Shadow Value of Funds

We work in the finance–binding region. Recall the Lagrangian

$$\mathcal{L}(k', \theta, \lambda; c, z) = U(k', \theta; z) + \lambda\big[c - k' M(k', \theta; z)\big],$$

where z is a primitive in $\{c, h, \tau, \pi, \rho, y\}$, $M(k', \theta; z)$ is given in (IA.4.2), and

$$S(k', \theta; z) \equiv M(k', \theta; z) + k' M_k(k', \theta; z).$$

[50] As noted previously, the equilibrium contractual rate r_c may also depend on borrower characteristics such as π and ρ, a dependence we suppress in the main text for simplicity. Allowing for this dependence would require additional mild regularity conditions for the π and ρ rows in Tables IA.4.1–IA.4.2, but would not affect the comparative statics with respect to R or the other primitives.

Table IA.4.2: Comparative statics inside the cap ($0 < \theta^* < \bar{\theta}$)

Primitive change	θ^*	$\lvert\varepsilon^{\text{int}}_{k',p_k}\rvert$	$\lvert\varepsilon^{\text{int}}_{k',R}\rvert$
Cash $c \downarrow$ (tighter liquidity)	$\uparrow$	$\uparrow$	$\uparrow$
Intangibles $h \downarrow$ (setup cheaper)	$\downarrow$	$\uparrow$	$\uparrow$
Install frictions $\tau \downarrow$ (easier)	$\downarrow$	$\uparrow$	$\uparrow$
Safer debt $\pi \uparrow$	$\downarrow$	$\downarrow$	$\uparrow$
Higher Recovery $\rho \uparrow$	$\downarrow$	$\downarrow$	$\uparrow$

Notes. Inside the cap, pledge reacts. Elasticities are

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

The first–order conditions in the interior ($0 < \theta^* < \bar{\theta}$) are

$$F_1(k',\theta,\lambda;z) \equiv U_k(k',\theta;z) - \lambda\,S(k',\theta;z) = 0, \tag{IA.4.7}$$

$$F_2(k',\theta,\lambda;z) \equiv U_\theta(k',\theta;z) - \lambda\,k'M_\theta(k',\theta;z) = 0, \tag{IA.4.8}$$

$$F_3(k',\theta;z,c) \equiv c - k'M(k',\theta;z) = 0. \tag{IA.4.9}$$

As before, the multiplier equals the marginal value of internal funds, $\lambda(c,z) = \partial V(c,z)/\partial c$.

Invoice-only tightness (fixed θ). For any fixed $\theta \in [0,\bar{\theta}]$, consider the reduced problem in (k',λ) with θ treated as a parameter. The FOCs are $F_1 = 0$ and $F_3 = 0$, and the derivative of λ with respect to a primitive z holding θ fixed is

$$\lambda^{\text{inv}}_z(k',\theta;z) = \frac{1}{S^2}\Big(SF_{1z} - k'M_zF_{1k}\Big), \tag{IA.4.10}$$

where $F_{1k} = U_{kk} - \lambda S_k$ and $F_{1z} = U_{kz} - \lambda S_z$, with $S_z = M_z + k'M_{kz}$.

We classify primitives as:

Definition 1 (Cost versus financial shifters)**.** *A primitive z is a* cost shifter *if it tightens the invoice and reduces marginal payoff:*

$$M_z > 0, \quad M_{kz} \geq 0, \quad U_{kz} \leq 0.$$

A primitive z is a financial shifter *if it relaxes the invoice and raises marginal payoff:*

$$M_z < 0, \quad M_{kz} \leq 0, \quad U_{kz} \geq 0.$$

In our baseline specification, h and τ are cost shifters, while π, ρ, and y are financial shifters.

Under the model assumptions, the proof in Appendix IA.4.6 for the cap case applies verbatim for any fixed θ and yields:

Lemma 1 (Invoice monotone tightness with fixed pledge)**.** *For any finance–binding optimum with θ treated as a parameter,*

$$\lambda^{\text{inv}}_c < 0, \qquad \lambda^{\text{inv}}_z > 0 \text{ for } z \in \{h,\tau\}, \qquad \lambda^{\text{inv}}_z < 0 \text{ for } z \in \{\pi,\rho,y\},$$

where λ^{inv}_z is defined in (IA.4.10).

Decomposition with endogenous pledge (interior region). When $0 < \theta^* < \bar{\theta}$, (k', θ^*, λ) solve the full system (IA.4.7)–(IA.4.9). Let

$$\lambda_z(c, z) \equiv \frac{\partial \lambda(c, z)}{\partial z}$$

denote the total derivative of λ with respect to z along this solution. By linearity of the implicit-function solution in the vector of right-hand sides, we can decompose

$$\lambda_z(c, z) = \lambda_z^{\text{inv}}\big(k'(c, z), \theta^*(c, z); z\big) \; + \; \lambda_z^{\theta}(c, z), \tag{IA.4.11}$$

where the first term is the invoice-only derivative (IA.4.10) evaluated at the interior solution $(k'(c, z), \theta^*(c, z))$, and the residual $\lambda_z^{\theta}(c, z)$ collects all effects operating through the endogenous adjustment of pledge $\theta^*(c, z)$. We refer to λ_z^{inv} as the *invoice channel* and to λ_z^{θ} as the *pledge-adjustment channel*.

Assumption 1 (Monotone tightness with endogenous pledge (MT-int))**.** *In any finance-binding interior equilibrium* $(k'(c, z), \theta^*(c, z), \lambda(c, z))$*:*

(a) For cost shifters $z \in \{h, \tau\}$*,*

$$\lambda_z^{\text{inv}}\big(k'(c, z), \theta^*(c, z); z\big) > 0 \quad \textit{and} \quad \lambda_z^{\theta}(c, z) \geq 0,$$

so that $\lambda_z(c, z) > 0$*.*

(b) For financial shifters $z \in \{\pi, \rho, y\}$*,*

$$\lambda_z^{\text{inv}}\big(k'(c, z), \theta^*(c, z); z\big) < 0 \quad \textit{and} \quad \lambda_z^{\theta}(c, z) \leq 0,$$

so that $\lambda_z(c, z) < 0$*.*

Equivalently, for each $z \in \{h, \tau, \pi, \rho, y\}$ *the pledge-adjustment channel does not overturn the sign of the invoice channel:*

$$\operatorname{sign}\big(\lambda_z(c, z)\big) = \operatorname{sign}\big(\lambda_z^{\text{inv}}(k'(c, z), \theta^*(c, z); z)\big).$$

Combining Lemma 1, Assumption 1, and the envelope property that $\lambda_c < 0$ in any finance-binding optimum yields:

Proposition 3 (Monotone tightness of λ)**.** *Under the model assumptions and 1, the shadow value of funds satisfies, in any finance-binding equilibrium (cap or interior),*

$$\frac{\partial \lambda}{\partial c} < 0, \quad \frac{\partial \lambda}{\partial h} > 0, \quad \frac{\partial \lambda}{\partial \tau} > 0, \quad \frac{\partial \lambda}{\partial z} < 0 \textit{ for } z \in \{\pi, \rho, y\}.$$

Lemma 2 (Interior pledge responses under MT-int)**.** *Suppose finance binds and* $0 < \theta^* < \bar{\theta}$*. Under* $m''(\theta) < 0$ *and Assumption 1, the optimal pledge satisfies*

$$\frac{\partial \theta^*}{\partial c} < 0, \quad \frac{\partial \theta^*}{\partial h} > 0, \quad \frac{\partial \theta^*}{\partial \tau} > 0, \quad \frac{\partial \theta^*}{\partial z} < 0 \textit{ for } z \in \{\pi, \rho, y\}.$$

Equivalently, in the direction of the comparative statics in Table IA.4.2:

$$c \downarrow \Rightarrow \theta^* \uparrow, \quad h \downarrow \Rightarrow \theta^* \downarrow, \quad \tau \downarrow \Rightarrow \theta^* \downarrow \; (\textit{weakly}), \quad z \uparrow \Rightarrow \theta^* \downarrow \textit{ for } z \in \{\pi, \rho, y\}.$$

Proof. From the interior θ–FOC $\lambda y m'(\theta^*) = \frac{(1-\pi)\rho}{R}(1 - m'(\theta^*))$ and $m''(\theta^*) < 0$ we have $\partial \theta^* / \partial \lambda > 0$ (recorded in Lemma IA.4.7). For any primitive z,

$$\frac{\partial \theta^*}{\partial z} = \frac{\partial \theta^*}{\partial \lambda} \frac{\partial \lambda}{\partial z},$$

so the signs of θ_z^* and λ_z coincide. The sign of λ_c is given by the envelope property in any finance-binding optimum; the signs of λ_h, λ_τ, λ_π, and λ_ρ follow from Proposition 3, which uses Lemma 1 and Assumption 1. □

IA.4.7 Lemmas on the shadow value λ

Let $\lambda \geq 0$ denote the common multiplier on the finance–binding constraints (budget and borrowing capacity), so that $\lambda = \mu$ when finance binds.

Lemma 3 (Interior mapping between λ and θ^*)**.** *If finance binds and the pledge choice is interior, $0 < \theta^* < \bar{\theta}$, then the interior pledge FOC (IA.4.3) implies*

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

and

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

Equivalently, $\partial\theta^*/\partial\lambda > 0$ in the interior region.

Lemma 4 (Cap KKT bound)**.** *If finance binds and the pledge cap binds ($\theta = \bar{\theta}$), the KKT condition associated with the upper bound implies*

$$\lambda\, y\, m'(\bar{\theta}) \;\geq\; \frac{(1-\pi)\rho}{R}\left(1 - m'(\bar{\theta})\right) \qquad \Longleftrightarrow \qquad \boxed{\lambda \geq \frac{(1-\pi)\rho}{R\,y}\,\frac{1-m'(\bar{\theta})}{m'(\bar{\theta})} \;=\; \frac{(1-\pi)\rho}{R\,y}\left(\frac{1}{m'(\bar{\theta})} - 1\right)}$$

IA.4.8 Proposition to diagnostic uniqueness

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 $1 + h/p_k - ym(\theta) > 0$, Assumption 1, and, at the cap for h, τ shocks, ($\kappa < 1$).*

Cap. *Define $S_{\text{cap}}(z) \equiv \left(\operatorname{sign}\Delta\lambda, \operatorname{sign}\Delta|\varepsilon_{k',p_k}|, \operatorname{sign}\Delta|\varepsilon_{k',R}|\right)$. 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 $S_{\text{int}}(z) \equiv \left(\operatorname{sign}\Delta\theta^*, \operatorname{sign}\Delta|\varepsilon_{k',p_k}|, \operatorname{sign}\Delta|\varepsilon_{k',R}|\right)$. 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 ((A5) 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 elasticities; (ii) $h\downarrow$ lowers θ^* while raising both elasticities; (iii) $\tau\downarrow$ slightly lowers θ^* while raising both elasticities; (iv) $y\uparrow$ lowers price sensitivity and raises rate sensitivity, with $\theta^*\downarrow$. Thus only $c\downarrow$ yields $(+,+,+)$. □